Developing EV Charge Station Health Indicators to Support Proactive Maintenance and Troubleshooting at GreenFlux

Author: Mohammad Wahyu Syafiul Mubarok, S6086853

Presented to: GreenFlux

Supervisors: Kris Borger, G.A.H. Laugs, Wided Medjroubi

Science, Business, & Policy

Faculty of Science and Engineering

University of Groningen

2026

GreenFlux Assets B.V.

Joan Muyskenweg 22
1096 CJ Amsterdam
The Netherlands

www.greenflux.com

Disclaimer

This report has been produced in the framework of an educational program at the University of Groningen, Netherlands, Faculty of Science and Engineering, Science Business and Policy (SBP) Curriculum. No rights may be claimed based on this report, other than described in the formal internship contract. Citations are only possible with explicit reference to the status of the report as a student internship product and written permission of the SBP staff.

Table of contents

Table of contents 3
Executive Summary 5
Part I: Introduction and Background 6
Part II: Analysis 9
2.1 External Environment Analysis: Reliability Pressure in the EV Charging Market 9
2.2 Internal Analysis: From Platform Capability to Uptime Intelligence 10
2.3 Competitor and Best-Practice Analysis: From CPMS Monitoring to Maintenance Intelligence 12
2.4 Current-State Operational Analysis 14
2.4.1 Understanding OCPP and Its Trends: Data Signals, Limitations, and Predictive-Maintenance Relevance 14
2.4.2 Current Troubleshooting Flow and Operational Pain Points 16
2.5 Data and Methodology 18
2.5.1 Data Structures 18
2.5.2 Methodology 20
2.6 Analytical Findings 21
2.6.1 Different Behavior Between AC and DC, and Why Predictive Maintenance is Different 22
2.6.2 Understanding Status Logs and Error Codes at Station and Connector Level 23
2.6.3 Alfen AC Representative Analysis 26
2.6.4 Alpitronic DC Representative Analysis 28
2.6.5 Summary of Analytical Findings 30
2.7 Strategic Synthesis and Discussion for GreenFlux 31
2.7.1 Synthesis of External, Internal, and Analytical Findings 32
2.7.2 SWOT and TOWS: Strategic Options for GreenFlux 33
2.7.3 Selected Strategic Direction: From Predictive Maintenance to Explainable Uptime Intelligence 34
Part III: Conclusion and Advice 35
3.1 Conclusion 35
3.2 Recommendations and Implementation Timeline 35
3.3 Feasibility, Organizational Readiness, and Adoption Risks 37
3.4 Reflection and Limitations 39

References 40

Appendix 44

Appendix 1. PESTEL Analysis: External Pressure on Charger Health and Uptime Intelligence 44

Appendix 2. Stakeholder Landscape and Power-Influence Analysis for Charger Reliability 48

Appendix 3. Internal Analysis: VRIO and 7S Assessment of GreenFlux's Maintenance-Intelligence Readiness 53

Appendix 4. Competitor and Best-Practice Analysis 58

Appendix 5. OCPP Signals and Troubleshooting Flows 63

Appendix 6. Data Structure and Methodological Details 71

Appendix 7. Detailed Analytical Evidence for Issue-Family Prioritization 80

Appendix 8. Implementation Support for Explainable Uptime Intelligence 103

Executive Summary

GreenFlux should not position predictive maintenance as an immediate fully automated Artificial Intelligence (AI) product. The recommended direction is to develop **explainable uptime intelligence**: a maintenance-intelligence layer that helps Charge Point Operator (CPO) customers prioritize which station or connector to review, understand the likely issue family, and choose the next troubleshooting route. This is the most credible bridge between current CPMS monitoring and future predictive alerts.

The business case is driven by charger network growth. As networks expand, CPOs face stronger pressure to maintain uptime, reduce troubleshooting delays, and manage larger volumes of operational data. GreenFlux already has rich CPMS/OCPP-related data, including status logs, session records, telemetry, vendor-specific fields, and metadata.

The project combined strategic assessment, competitor review, stakeholder workflow insights, and AC/DC charger data analysis, using Alfen as the representative AC case and Alpitronic as the representative DC case. The analysis shows that maintenance intelligence should not rely on raw error-code counts or one generic charger-health score. AC and DC chargers behave differently, station and connector logs are not interchangeable, and broad categories such as Other Error become useful only when combined with Vendor Error Code, Info, recurrence, entity level, and session context.

The strongest first candidates are AC metering/P1/Modbus/measurement, AC grounding/electrical protection, and DC stack/power-module availability. These issue families provide the most credible starting point because they show stronger evidence for operational interpretation and review prioritization. Other DC session-flow issues may also be useful, but should be treated with more cautious model-quality labels.

The implementation should be staged. GreenFlux should first define and dashboard selected high-evidence issue families, then add station-level and connector-level review queues with reason-code explanations, model-quality labels, and troubleshooting routes. Before moving toward automated alerts, GreenFlux should build feedback loops linking ranked cases to support tickets, remote reset/reboot outcomes, vendor diagnosis, technician findings, false positives, and confirmed repairs.

Organizationally, maintenance intelligence should be treated as a cross-functional product capability, not only as a data-science model. GreenFlux should assign a clear internal product owner, supported by product, data/engineering, support, customer success, and CPO-facing teams. The main adoption risks are false confidence, false positives, workflow mismatch, vendor-specific behavior, and insufficient feedback. These risks can be reduced by starting with narrow pilots, keeping humans in the loop, preserving AC/DC and station/connector context, and using explainable outputs rather than black-box scores.

The main limitation is that the analysis relies on issue-family indicators rather than confirmed physical root-cause labels; recurrence may also indicate unresolved issues rather than true early warning. Overall, GreenFlux should pursue predictive maintenance as a longer-term capability, but its immediate opportunity is explainable uptime intelligence: an evidence-based, operationally useful, commercially relevant, and feasible path toward more reliable charging operations.

Part I: Introduction and Background

The energy transition is not only about replacing fossil-based technologies with low-carbon alternatives. It also depends on whether those alternatives can work reliably in everyday life. Transport electrification shows this clearly. Global electric car sales exceeded 20 million in 2025 and reached around 25% of total car sales, while charging infrastructure continued to expand in parallel (International Energy Agency [IEA], 2026). In 2025, the global stock of public charging points reached more than 7 million, while Europe public charging points increased by around 20%. The Netherlands had the highest number of public charging points in Europe by the end of the year, with around 210,000 points, followed by Germany and France (IEA, 2026). Electric Vehicle (EV) charging has therefore become more than an infrastructure-deployment issue. It is now an operational layer of the energy transition.

As charging networks grow, the question is no longer only how many charge points are installed. The more practical question is whether they are available, usable, and dependable when drivers need them. Public charging reliability affects user confidence, Charge Point Operator (CPO) revenue, asset utilization, support workload, and the credibility of transport electrification. Recent scientific literature supports these shifts from deployment quantity to operational reliability. Singh et al. (2025) show that poor perceived public-charging reliability can reduce EV purchase intention, meaning that unreliable charging can slow market growth even when infrastructure is being expanded (Singh et al., 2025). Equity-focused research also shows that public charging is not only a question of spatial access, but also of reliability, usability, cost burden, awareness, and trust (Chen et al., 2025). Previous work on public charging user experience similarly identifies availability, reliability, and interoperability as key components of a functional charging ecosystem (Bernard, 2023; Bernard et al., 2024). In this sense, charger reliability matters both as an operational performance issue and as part of the wider conditions for credible and equitable electrification.

Reliability is also becoming more visible in regulation and standardization. In the United Kingdom, the Public Charge Point Regulations 2023 guidance requires rapid public charge point networks to achieve 99% annual reliability, measured across a CPO's rapid network (Department for Transport, 2024). In the United States, the National Electric Vehicle Infrastructure (NEVI) standard requires each charging port to have an average annual uptime greater than 97% (Federal Highway Administration, 2023). The Open Charge Alliance also notes tat Open Charge Point Protocol (OCPP)-based charger-management messages and configuration options can support downtime detection and, in some cases, downtime prevention (Open Charge Alliance [OCA], 2024). These examples indicate a broader direction: uptime is becoming a matter of service quality, performance monitoring, and operational accountability, not only a back-office maintenance concern.

However, uptime alone is not sufficient to guide maintenance and troubleshooting. A charger may be technically online but still unreliable from an operational or user-experience perspective, for example because of recurring communication instability, connector-level faults, failed charging session, inaccurate status reporting, or repeated recoverable errors. Recent work on charging-site performance argues that conventional uptime metrics can be too aggregated to support operational improvement because they often lack diagnostic granularity for site

operators (D'Agostino et al., 2026). Empirical studies also show that public-charging reliability problems can be hidden behind reported availability figures and that clearer definitions of reliability, uptime, downtime, and excluded time are needed for public charging infrastructure (Rempel et al., 2024; Webb et al., 2023). Therefore, the operational challenge is to understand what type of issue is occurring, where it appears, and whether it requires review.

GreenFlux is positioned within this reliability challenge as an EV charging software and charge point management platform provider. Its platform helps CPOs and eMobility providers operate charging businesses at scale, with capabilities related to charge point management, monitoring, smart charging, payments, roaming, and operational control (GreenFlux, n.d.). This position is important because GreenFlux operates on the software and data layer between charging hardware, CPO operations, support processes, and end-user charging experience. It does not manufacture chargers or perform all field maintenance itself, but its platform receives operational signals from connected charging assets. This role is consistent with the broader function of charging management systems, which can support remote operation, diagnostics, error reporting, maintenance scheduling, monitoring, control, data processing, and decision-making (Hsaini et al., 2022; Martins & Rodrigues, 2025; Ozkaya et al., 2026). This report therefore treats EV charger reliability not only as a hardware-maintenance problem, but as a digital-operational intelligence challenge in which charge points management software (CPMS) platforms can translate fragmented charger data into actionable maintenance decisions.

This creates both an opportunity and a limitation. Charging stations continuously generate operational data, including status notification, transaction records, error codes, vendor-specific messages and other platform signals. In principle, these data can support earlier detection of deteriorating charger conditions and more efficient troubleshooting. Research on EV charging big data shows that operational data can be used to identify features related to the health status of charging devices, support fault diagnosis models, and improve user experience (Wu et al., 2025; Sakwa et al., 2024). However, raw charger data does not automatically become maintenance intelligence. The signals are often noisy, repeated, vendor-specific, and shaped by differences in protocol version, charger type, and operational context.

The central problem is therefore not a lack of data, but the difficulty of converting fragmented operational logs into meaningful signals for maintenance and troubleshooting. For GreenFlux and its CPO customers, the challenge is to distinguish true failure from noise, prioritize issues, understand likely causes, and support faster resolution. This requires more than counting errors or detecting when a charger becomes unavailable. It requires a structured interpretation of charger behavior, failure context, and support workflow.

While existing studies increasingly examine charging reliability, uptime, intelligent monitoring, and predictive maintenance, less attention has been given to how CPMS providers can operationalize fragmented charger logs into explainable prioritization and troubleshooting support for CPOs. It also requires caution in how model outputs are used. Predictive maintenance systems are more likely to be adopted when they remain explainable, human-in-the-loop, and aligned with maintenance workflows, rather than operating as opaque automated decision systems (Amaliah et al., 2025; Cummins et al., 2024).

Several concepts guide this report. Charger uptime refers to whether charging infrastructure is available and functional for intended use, while incident management covers the detection, triage, diagnosis, and resolution of charging problems. Unhealthy EV charging is used as a data-driven proxy for abnormal charger-status or charging-session patterns, such as degraded performance, interrupted sessions, repeated fault-related signals, or conditions requiring operational attention. It is not treated as confirmed charger failure, since irregularities may also originate from user behavior, vehicle-side issues, communication problems, grid or site conditions, or temporary operational states. Charger health and uptime intelligence refers to the capability of turning raw charger data into actionable insights, diagnostic guidance, and workflows that help CPOs maintain reliable charging services. Predictive maintenance is therefore treated as a more advanced stage of this capability, not as the immediate starting point.

Against this background, the main research question of this report is **How can GreenFlux best enable its CPO customers to efficiently maintain and troubleshoot their charge stations?**

To answer this question, the study analyzes available charging-station and platform data and translates the findings into strategic advice for GreenFlux. Predictive maintenance is treated as a staged capability, built first through signal interpretation, issue prioritization, diagnostic support, and operational workflows before moving toward more autonomous models. Key insights are presented in the main body, with technical details and supporting analyses provided in the appendix.

Reading Guide

This report is structured as an advisory report that moves from context to analysis and then to implementation advice. **Part 1** introduces the relevance of charger reliability within the energy transition, explains GreenFlux's position in the EV charging ecosystem, and defines the concepts used throughout the report.

Part 2 presents the main analysis. It examines the external environment through PESTEL, assesses GreenFlux's internal capabilities using VRIO, and compares competitor and best-practice developments in CPMS-based reliability features. It then explains the operational and data context, including OCPP, status notifications, error codes, and troubleshooting flows. The methodology section summarizes the analytical approach, while detailed method exploration is placed in the appendix. The findings compare AC and DC charger behavior, station- and connector-level signals, hardware-oriented errors, and representative Alfen and Alpitronic cases.

Part 3 translates the findings into advice, including the conclusion, recommendations, implementation timeline, stakeholder value, feasibility, and adoption risks. The appendix provides supporting analyses behind the main insights.

Part II: Analysis

2.1 External Environment Analysis: Reliability Pressure in the EV Charging Market

The external environment indicates that charger reliability is becoming a strategic issue for charge point operators and charge point management platform providers. A PESTEL analysis was used to identify the policy, market, societal, technological, environmental, and legal pressures most relevant to charger uptime and maintenance intelligence (Whittington et al., 2023). The purpose of this section is not to provide a full overview of the EV charging market, but to identify the external forces that shape why GreenFlux should treat charger health and uptime as a product and service opportunity. The full PESTEL analysis is provided in **Appendix 1**, while full stakeholder landscape and power-influence analysis for charger reliability is shown in **Appendix 2**.

The first pressure comes from regulation and public policy. EV charging policy is moving beyond infrastructure deployment toward infrastructure quality, interoperability, user information, payment accessibility, and service reliability. In the European Union, the Alternative Fuels Infrastructure Regulation has been applied since April 2024 and aims to ensure minimum infrastructure, full interoperability, comprehensive user information, and adequate payment options for alternative-fuel infrastructure (European Commission, 2024). In the United Kingdom, rapid public charge point networks are required to achieve 99% annual reliability across each CPO's rapid network (Department for Transport, 2024). In the United States, the NEVI standard requires each charging port to have an average annual uptime greater than 97% (Federal Highway Administration, 2023). Although these examples come from different regulatory contexts, they point in the same direction. Charging reliability is becoming a measurable performance expectation, not only a voluntary service-quality ambition.

The second pressure is economic and operational. As charging networks scale, CPOs need to protect asset utilization, reduce avoidable downtime, and manage maintenance costs more efficiently. A charger that is installed but frequently unavailable creates limited value for the operator and can increase customer-support workload. This makes uptime not only a technical metric, but also a business-performance issue. For CPMS providers such as GreenFlux, this changes customer expectations. Basic connectivity and status visibility remain necessary, but they are not sufficient if CPOs need faster triage, clearer fault interpretation, and better prioritization of maintenance actions.

The third pressure is technological. Charging infrastructure is becoming more complex because networks include different charger types, vendors, firmware versions, OCPP versions, and AC/DC configurations. OCPP provides a communication layer between chargers and management systems, but it does not automatically create diagnostic clarity. The Open Charge Alliance notes that OCPP itself is only a set of messages between a charger and its management system (Open Charge Alliance [OCA], 2024). This means that the strategic value for GreenFlux lies not only in receiving charger data, but in interpreting fragmented messages into operationally meaningful signals.

The fourth pressure is social and energy transition. Public charging reliability influences user confidence in electric mobility. User-experience research identifies availability, reliability, and interoperability as key dimensions of public charging usability (Bernard et al., 2024). Reliability also has an equity dimension. Evidence from the United States suggests that disadvantaged communities face fewer public charging stations and more reliability issues, meaning that unreliable charging can create unequal barriers for users with fewer alternatives (Yu et al., 2025). From an energy-transition perspective, reliable charging infrastructure is therefore part of the enabling system that makes low-emission mobility usable in practice.

Stakeholder power and influence shape how these external pressures translate into GreenFlux's strategic opportunity. CPO customers are the primary high-power stakeholders because they are responsible for charger uptime, customer experience, operational cost, and service accountability. Regulators and public authorities also have high influence because they define uptime and service quality expectations. Charger OEMs, vendor diagnostic systems, field-service partners, site hosts, and grid or facility operators influence whether issues can be interpreted, escalated, and resolved. EV drivers have less direct technical control, but their charging experience affects trust, demand, and reputational pressure on CPOs. This means GreenFlux operates in a multi-actor reliability landscape. It cannot directly solve every hardware, site, or vehicle-side problem, but it can improve how operational signals are interpreted, prioritized, and routed across stakeholders.

Taken together, the external environment suggests that GreenFlux operates in a market where uptime is becoming more visible, more measurable, and more strategically important. Regulation increases accountability, CPO profitability increases the need for operational efficiency, charger technology creates richer but more complex data, stakeholder coordination shapes implementation, and EV users increasingly expect charging services to be dependable.

External implication for GreenFlux

The external environment shows a shift from "more chargers" toward "more reliable charging services." For GreenFlux, this creates an opportunity to move beyond basic charge point monitoring toward charger health and uptime intelligence. Rather than replacing CPOs, vendors, or field-service actors, this capability can help connect operational data with clearer monitoring, earlier issue awareness, and better coordination across the charging ecosystem.

2.2 Internal Analysis: From Platform Capability to Uptime Intelligence

The external analysis shows that charger reliability is becoming more visible, measurable, and strategically important. The internal question is whether GreenFlux has the resources and capabilities to respond to this shift. A qualitative VRIO analysis was used to assess GreenFlux's internal position in relation to charger health and uptime intelligence. The framework is grounded in the resource-based view of the firm, which links competitive advantage to firm-specific resources and capabilities (Barney, 1991), and is commonly used in strategic management to assess whether organizational capabilities can contribute to advantage (Whittington et al., 2023). A compact 7S assessment is then used to examine whether the organizational conditions around

strategy, structure, systems, skills, staff, style, and shared values are aligned with this direction. The 7S framework is relevant because it highlights that organizational capability depends not only on formal structure, but also on the alignment of hard and soft organizational elements (Waterman et al., 1980). The full VRIO and 7S assessments are provided in **Appendix 3**.

GreenFlux has a strong starting position because it operates on the software and data layer of the EV charging ecosystem. Its charge point management platforms, APIs, monitoring functions, roaming operations, smart charging capabilities, and support-related services give the company experience in managing complex charging operations at scale. These capabilities are valuable because they place GreenFlux close to the operational signals that matter for uptime, including status notifications, transaction records, error codes, vendor-specific messages, and platform-side operational information. They also give GreenFlux the possibility to embed new reliability features into existing CPO workflows, rather than building separate tools outside the platform.

At the same time, this foundation should not be overstated. The VRIO assessment suggests that GreenFlux's advantage does not come from platform data access alone. Many CPMS providers can receive charger status and transaction data through OCPP-based communication, and basic monitoring or error visibility is increasingly part of the expected CPMS feature set. The main question is whether GreenFlux can organize fragmented charger data into signals that are meaningful for operations. This includes structuring station and connector-level events, distinguishing repeated noise from relevant patterns, interpreting vendor-specific error messages, linking technical events to charging-session outcomes, and translating these insights into practical triage and troubleshooting guidance.

GreenFlux also has adjacent strengths that can support this direction. Its smart charging and energy-management algorithms show experience with algorithm-supported charging operations, while its interoperability, roaming, and Roaming-as-a-services (RaaS) activities show that the company can operate across multi-actor and multi-protocol charging ecosystems. These capabilities strengthen GreenFlux's credibility as a platform operator, but they should not be confused with a proven predictive-maintenance capability. Predictive maintenance requires different evidence bases, such as validated failure signals, reliable incident labels, feedback from support or field maintenance, and integration into daily operational decision-making.

Beyond resources and platform capabilities, the 7S assessment indicates that the main internal readiness issue is organizational alignment. Maintenance intelligence is strategically consistent with GreenFlux's CPMS role, but it is not yet established as a mature core product layer. Current ownership is distributed across product, data, support, and customer-facing functions, while product teams would likely hold the development decision. The main gap is therefore not only technical. GreenFlux also needs reusable feature tables, structured review outcomes, feedback labels, and workflow integration. It also requires hybrid skills that combine OCPP knowledge, charger diagnostics, data/Machine Learning capability, product thinking, and support-process understanding. This suggests that GreenFlux has the strategic fit and relevant capabilities, but turning into maintenance intelligence into a scalable feature will require clearer cross-functional ownership and feedback loop.

The most relevant internal opportunity is therefore not data access itself, but the operational judgement built around it. GreenFlux already has management and monitoring capabilities, but this area is likely closer to a temporary advantage than a sustained moat unless it can demonstrate distinctive interpretation, prioritization, and workflow value. In practice, this means developing vendor-aware error interpretation, connector-level diagnostics, alert prioritization, health indicators, and review workflows that CPOs can trust.

The internal analysis therefore points to a staged opportunity. GreenFlux is well positioned to develop charger health and uptime intelligence, but the strongest starting point is not fully automated prediction. A more feasible path is to first improve signal interpretation, diagnostic support, and issue prioritization, then use structured operational feedback to strengthen health indicators and predictive-maintenance readiness. Strategically, GreenFlux's potential advantage will depend on whether it can combine platform data, analytics, charger knowledge, workflow design, and cross-functional ownership into features that help CPOs act faster and more consistently.

Internal implication for GreenFlux

GreenFlux has the right platform position, but the defensible advantage still needs to be built. The opportunity lies in converting fragmented OCPP, vendor, station, connector, and session signals into maintenance support that CPO customers can use. To make this scalable, GreenFlux will need clearer cross-functional ownership, structured feedback labels, and workflow integration between product, data, support, and customer-facing teams.

2.3 Competitor and Best-Practice Analysis: From CPMS Monitoring to Maintenance Intelligence

The competitor and best-practice analysis was conducted to understand how the CPMS market is positioning reliability, diagnostics, and predictive maintenance features. The purpose is to assess whether charger health and uptime intelligence are becoming relevant areas of differentiation. Porter's Five Force was used to interpret the competitive structure of the market, while a benchmark analysis was used to compare selected CPMS providers and adjacent maintenance-intelligence players. The detailed benchmark and Porter analysis are provided in **Appendix 4**.

The CPMS market already contains many providers with overlapping baseline capabilities. Mature platforms commonly position themselves around charge point monitoring, billing, payments, roaming white-label applications, APIs, energy management, and operational control. Road.io, for example, positions itself as a modular white-label EV charging platform that manages payments, roaming integrations, compliance, and driver support (Road, n.d.). Chargecloud describes CPMS as a system for digitally managing charging networks, including communication with stations, connector entity, pricing, invoicing, and remote troubleshooting of hardware faults (Chargecloud, 2025). ChargePoint similarly presents its software platform around station management, proactive monitoring, diagnostics, interoperability, and load management (ChargePoint, 2025). These examples suggest that basic CPMS functionality is increasingly becoming a competitive baseline rather than a clear source of differentiation.

The more relevant competitive shift is the movement from monitoring toward maintenance intelligence. Public product materials show three broad patterns. First, mature CPMS providers increasingly add proactive alerts, remote recovery, and issue-management features. AMPECO, for example, describes real-time issue detection, automated fault-recovery algorithms, and an issue-management flow from detection to resolution (AMPECO, 2026). Second, some platforms use stronger language around self-healing and automated remediation. Monta's alert feature states that it automatically fixes what it can and alerts the right people when it cannot, while Virta claims that its self-healing functionality can automatically solve up to 30% of charger issues (Monta, 2025; Virta, 2024). Third, specialist or adjacent players are positioning Artificial Intelligence (AI) diagnostics and predictive maintenance as their core offer. EVailable.ai presents an OCPP-based solution for analytics, predictive maintenance, and automation, while Ampcontrol frames 24/7 monitoring and predictive maintenance as ways to improve uptime and reduce repair costs (Ampcontrol, 2024; EVailable.ai, n.d.).

The benchmark indicates that charger reliability is becoming a more visible part of CPMS differentiation. However, it also shows that public claims vary in specificity and evidence strength. Since this benchmark is based on public product materials, it reflects market positioning rather than independently verified model performance or tested maintenance outcomes. Terms such as "AI," "predictive," "self-healing," and "automatic diagnostics" are increasingly common, but these labels should be interpreted cautiously. Predictive-maintenance literature stresses that model outputs need to be explainable, validated, and usable by human operators, especially when they support operational decisions (Cummins et al., 2024; Amaliah et al., 2025). Similarly, trustworthy AI guidance emphasizes validity, reliability, transparency, explainability, and accountability as key conditions for responsible AI use (NIST, 2023). A feature that detects an issue is not the same as one that diagnoses the likely cause, prioritizes operational urgency, supports remote recovery, or predicts failure using validated evidence.

Porter's Five Forces reinforce that GreenFlux operates in a competitive CPMS market where differentiation through generic platform functionality is increasingly difficult. Rivalry is high because many providers offer similar combinations of monitoring, billing, roaming, energy management, and operational features. Buyer power is also high, as sophisticated CPOs, fleets, utilities, and enterprise customers can compare platforms based on uptime support, multi-OEM integration, compliance readiness, and operational efficiency. Substitution pressure is moderate to high because AI diagnostics firms, OEM tools, uptime specialists, and internal CPO analytics teams can address parts of the maintenance problem outside a traditional CPMS. Entry barriers are moderate: software entry is possible, but enterprise-scale CPMS delivery requires integration capability, certifications, operational credibility, customer trust, and support capacity. Supplier power is also moderate because CPMS performance still depends on charger OEM behavior, firmware stability, protocol quality, and OCPP implementation consistency. This reflects Porter's view that competition is shaped not only by direct rivals, but also by buyers, suppliers, substitutes, and potential entrants (Porter, 2008).

The strategic implication is that GreenFlux should not compete only on generic CPMS functionality or broad "AI prediction" claims. A more defensible opportunity lies in explainable operational judgement: helping CPOs understand which asset is affected, whether the signal appears at station or connector level, how urgent it is, and what action should follow. These points toward vendor-aware classification, alert prioritization,

diagnostic guidance, and health indicators validated against charging behavior, allowing GreenFlux to move from basic monitoring toward charger health and uptime intelligence without overclaiming what current data can predict.

Competitive implication for GreenFlux

The market is moving from CPMS connectivity toward operational reliability tooling. GreenFlux's opportunity is not to copy broad AI-predictive language, but to build a more explainable and workflow integrated reliability layer around the charger data it already receives.

2.4 Current-State Operational Analysis

The previous sections show that charger reliability is becoming strategically important, that GreenFlux has a relevant platform position, and that the CPMS market is moving toward maintenance intelligence. However, before proposing data-driven maintenance features, it is necessary to understand the current operational reality. What GreenFlux can observe through platform data, what the data can and cannot explain, and why troubleshooting still requires operational interpretation. This section therefore examines the role of OCPP-based charger data and the troubleshooting flow through which charging issues are currently detected, interpreted, routed, and resolved.

The analysis is divided into two parts. **Section 2.4.1** explains the OCPP and CPMS data reality, focusing on the available signals, their diagnostic value, and their limitations for predictive-maintenance readiness. **Section 2.4.2** then uses interview-based troubleshooting flows from CPMS, DC charger owner/operator, and AC charger owner/operator perspectives to identify operational pain points. Detailed OCPP data elements and interview-based flows are provided in **Appendix 5**.

2.4.1 Understanding OCPP and Its Trends: Data Signals, Limitations, and Predictive-Maintenance Relevance

Open Charge Point Protocol (OCPP) is the main communication layer between charge points and charge point management systems. It allows charging stations to exchange operational messages with a backend system, making it possible for a CPMS such as GreenFlux to monitor charger status, transaction, connector states, error codes, and other operational events. The Open Charge Alliance describes OCPP as a global open communication protocol between charging stations and charging management systems, supporting interoperability and reducing dependence on proprietary systems (Open Charge Alliance [OCA], 2024). In this project, OCPP matters because it defines much of the raw operational data that can be used for charger monitoring, troubleshooting, and future maintenance intelligence.

The value of OCPP data lies in operational visibility. Status notifications can show whether a charger or connector is available, unavailable, faulted, charging, suspended, or recovering. Error codes provide a first-level indication of abnormal conditions. Vendor error codes and information fields may add more specific context. Transaction or session records can show whether a charging attempt started, stopped, delivered energy, or ended unusually. Together, these signals create a foundation for understanding charger behavior across large networks. They also make it possible to identify repeated faults,

recurring abnormal sessions, charger-type differences, and station or connector-specific patterns.

However, OCPP data should not be treated as a complete diagnostic system by itself. A status notification or error code shows that something was reported, but it does not always explain the underlying cause. In this report, OCPP messages are treated as operational signals, not direct diagnoses. They become useful maintenance indicators only after being structured, contextualized, and linked to session outcomes or operational evidence. The meaning of signal depends on the OCPP version, charger vendor, firmware version, charger type, station-versus-connector reporting level, and surrounding events sequence. This is particularly important for broad error categories such as Other Error, where the operational meaning may only become visible after checking vendor-specific fields, information messages, and the sequence of events around the charging session.

Error-code consistency is another important limitation. The ChargeX Consortium argues that standardized minimum error codes are needed to simplify troubleshooting and improve charging reliability, because inconsistent error reporting makes diagnosis and resolution harder across the EV charging industry (Savargaonkar et al., 2023). This is directly relevant to GreenFlux because a CPMS provider operates across heterogenous charger fleets. If different vendors use different error semantics, the platform cannot rely only on raw error labels. It needs a translation layer that groups signal into meaningful issue families and link them to likely operational actions.

Charger availability and charging success are not the same. A charger may appear technically available but still fail to deliver successful user experience because of authorization, payment, communication, connector, vehicle-side, or power-transfer problems. The EV-WATTS reliability study uses charging-session data, error codes, and power-transfer data to quantify reliability from the driver-experience perspective, showing that session outcomes and delivered power are important complements to status information (CLEAResult, 2024). These supports link OCPP status logs with transaction and session data rather than interpreting status notifications in isolation.

Protocol development should be interpreted pragmatically. This analysis relies mainly on OCPP 1.6-style signals, including status notifications, error codes, vendor-specific information, session records, and station- or connector-level identifiers. These signals are imperfect, but they can support better issue interpretation when structured and connected to session outcomes. OCPP 2.x points to richer future diagnostics: OCA's uptime-monitoring material notes that OCPP 2.0.1 can use monitoring support and device-model variables to detect threshold exceedances, enabling more granular remote monitoring than simple status updates alone (OCA, 2024). The implication is that GreenFlux should improve the interpretation of current OCPP 1.6 data while preparing for more detailed OCPP 2.x diagnostics.

Predictive-maintenance readiness cannot start from raw errors alone. OCPP messages need operational context: asset level, station-versus-connector reporting, recurrence, vendor-specific meaning, and links to failed or abnormal sessions. OCPP provides the data foundation, but GreenFlux's value lies in converting that foundation into operational judgement.

Data implication for GreenFlux

OCPP gives GreenFlux visibility into charger operations, but visibility is not the same as diagnosis. The opportunity is to transform raw OCPP, vendor, station, connector, and session signals into structured maintenance indicators that can support troubleshooting, prioritization, and future predictive-maintenance readiness.

2.4.2 Current Troubleshooting Flow and Operational Pain Points

While OCPP explains what GreenFlux can observe, it does not explain how those signals are used in real troubleshooting. To understand this operational layer, troubleshooting was mapped from three perspectives: CPMS/platform perspective, DC charger owner/operator perspective, and AC charger owner/operator perspective. These flows should be read as interview-based operational mappings rather than exhaustive process audits. The purpose is to identify recurring troubleshooting bottlenecks across stakeholder perspectives, not to claim that all CPMS, DC, or AC operators follow the exact same process.

Across the three perspectives, the troubleshooting process follows a broadly similar structure as can be seen in **figure 1**. Starting with an issue becomes visible, the platform or operator checks available information, a basic remote recovery action may be attempted, the issue is interpreted, and the case is either resolved remotely or escalated to customer service, vendor support, an installation engineer, or a field technician. This is consistent with wider reliability literature that treats charging failures as a process chain problem. A failed or poor charging experience may originate from authorization, communication, payment, charger status, user action, power transfer, or maintenance response, rather than from charger hardware alone (Roeckle et al., 2022; CLEAResult, 2024). The detailed interview-based flows are provided in **Appendix 5**.

The three interview-based flows show different versions of the same operational problem: issue signals are available, but their operational meaning is not always clear. From the CPMS perspective, a user-facing problem may occur even when the diagnostic log shows *No Error*, while failed sessions may originate from authorization, token/card use, driver behavior, or communication problems rather than charger failure. From the DC owner/operator perspective, reset or reboot actions may solve a meaningful share of issues, but unresolved cases still require false-positive filtering, historical error review, site validation, vendor guidance, or technician escalation. From the AC owner/operator perspective, ambiguous errors such as *Other Error* and mixed error combinations can make it difficult to decide whether the issue is installation-related, hardware-related, communication-related, or user-context related. Across these perspectives, the recurring pain point is not the absence of data, but the lack of a systematic layer that turns issue signals into operational judgement: what happened, where it happened, whether it matters, and what action should follow.

This current-state analysis also exposes a weak feedback loop. Field intervention, vendor support, or installation fixes may resolve the issue, but confirmed root cause is not always linked back to the original CPMS log pattern. This limits predictive-maintenance readiness because models or rule-based indicators need validated outcomes to distinguish meaningful warning signals from noise. The ChargeX minimum required error code work supports this issue at an industry level: inconsistent error

reporting reduces the ability of operators to interpret, diagnose, and resolve charging problems consistently across vendors and systems (Savargaonkar et al., 2023).

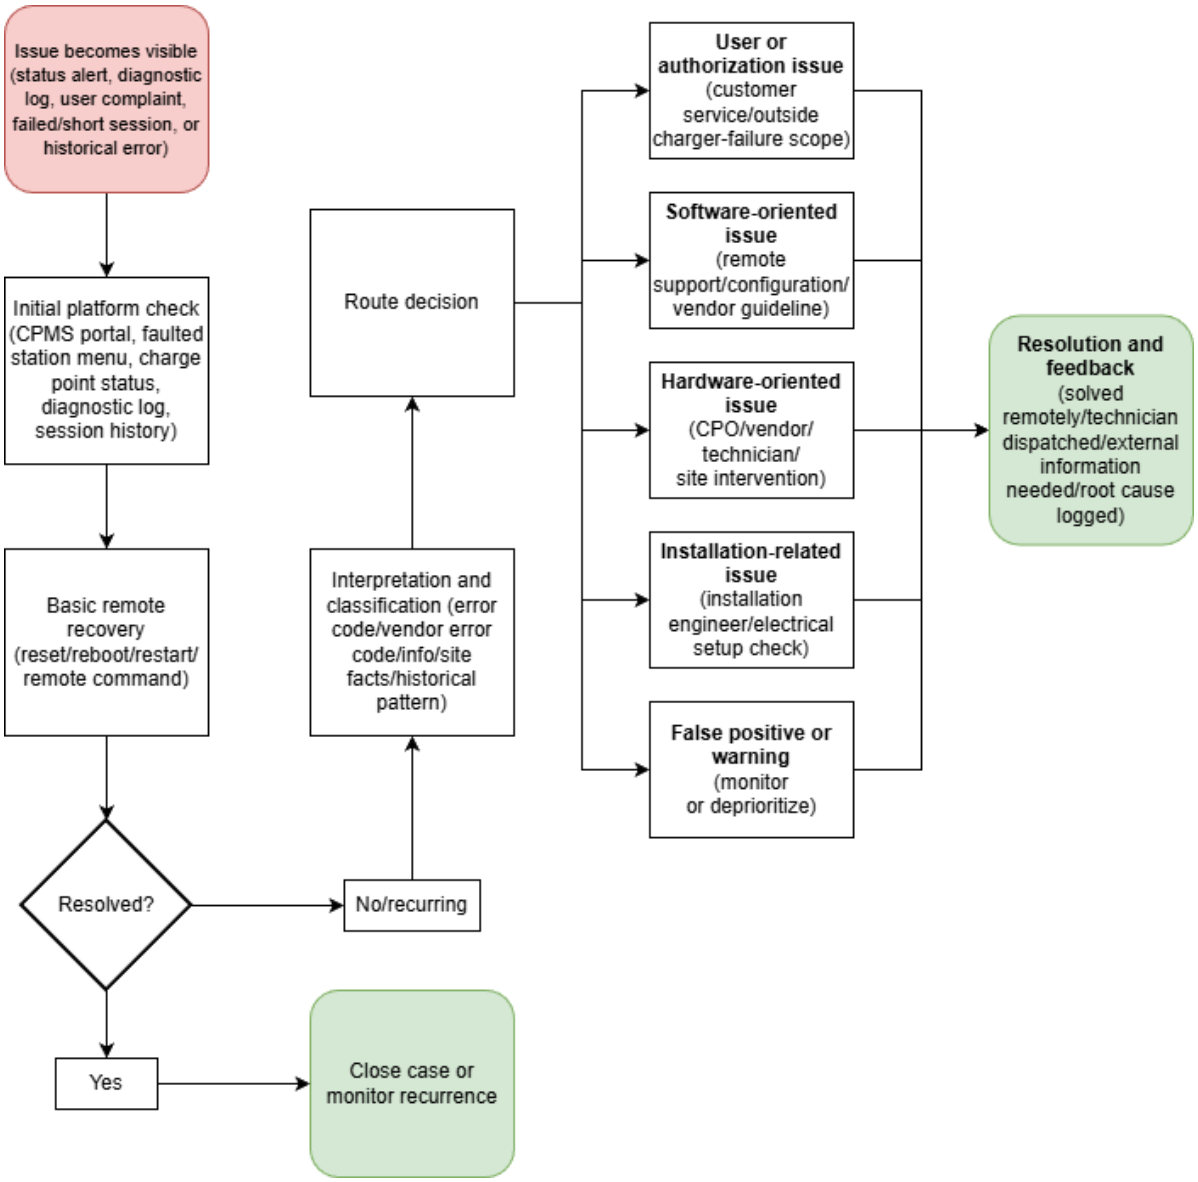

Figure 1. Synthesized troubleshooting flow based on CPMS, DC charger owner/operator, and AC charger owner/operator interviews.

For GreenFlux, the implication is clear. The current troubleshooting challenge is not only to detect charger issues, but to interpret, classify, prioritize, and route them. This supports the analytical choices in the next section, by separating station and connector-level signals, compressing repeated events, distinguishing AC and DC behavior, linking events to session outcomes, and treating unhealthy charging as a data-driven proxy rather than confirmed root-cause diagnosis.

Operational implication for GreenFlux

GreenFlux can create value by turning raw OCPP, vendor, station, connector, and session signals into structured operational context. The strongest near-term opportunity is not full failure prediction, but better triage: what happened, where it happened, whether it matters, and what action should follow.

2.5 Data and Methodology

This section summarizes the selected methodology, while detailed data preparation, method exploration, screening logic, and modeling setup are provided in **Appendix 6**. The methodological challenge is that the data contains rich operational signals but limited verified root-cause or repair-outcome labels. For this reason, the analysis does not start from the complex black-box prediction. Instead, it prioritizes interpretable data structuring, issue-family mapping, feature engineering, and baseline issue-likelihood ranking.

Predictive-maintenance literature includes a wide range of approaches, including rule-based monitoring, supervised machine learning, anomaly detection, survival or remaining-useful-life modeling, deep learning, and explainable Artificial Intelligence (AI). However, method suitability depends on data structure, label maturity, temporal continuity, interpretability needs, and deployment context. Reviews of machine-learning applications in predictive maintenance emphasizes that model performance depends not only on algorithm choice, but also on data quality, feature construction, and meaningful failure definition (Carvalho et al., 2019; Çinar et al., 2020). Therefore, the selected methodology is intentionally conservative but not technically narrows. It reflects the current maturity of the available GreenFlux data and the operational need for explainable decision support.

The methodology follows staged logic. First, raw OCPP/status, charging-session, telemetry, and metadata records were converted into station-day and connector-day tables/panels. Second, raw status and error signals were screened to distinguish high-frequency events from potentially fault-linked or operational relevant signals. Third, Error Code, Vendor Error Code, and Info fields were mapped into interpretable issue families. Fourth, candidate features were engineered around recurrence, session quality, exposure, telemetry, metadata, and diagnostic context. Finally, baseline logistic-regression models were used to test whether the selected features could support explainable issue-likelihood ranking outside the training period.

The selected methodology should be interpreted as a bridge between current monitoring and future predictive maintenance. It tests whether GreenFlux can move from raw log inspection toward structured issue-family ranking, while leaving calibrated predictive alerts and more complex models for later validation. This distinction matters because interpretable models are often preferable when operational users need to understand why a case is prioritized, especially in decision-support contexts where black-box explanations may create false confidence (Rudin, 2019).

2.5.1 Data Structures

The analysis used several GreenFlux data layers: charger metadata, station-level status notifications, connector-level status notifications, charging-session records, Meter Values/telemetry where available, and vendor-specific diagnostic fields such as Vendor Error Code and Info. These data sources have different structures. Status notifications are event-based, CDRs are session-based, Meter Values are high-frequency telemetry records, and charger metadata is mostly static. This makes direct comparison difficult without first creating a common analytical unit. In addition, the period of dataset are three months: train dataset (01/10/2025 – 30/11/2025); test dataset (01/12/2025 – 15/12/2025); and keep-out dataset (16/12/2025 – 31/12/2025). The charging type data

are business and public, because it needs more attention on uptime to generate more revenue on CPO side.

To align these different structures, the analysis converted raw data into station-day and connector-day panels. A station-day represents one charge station on one date, while a connector-day represents one connector entity on one date. This allows status events, session outcomes, telemetry, metadata, and recurrence features to be compared within a consistent operational unit. It also supports train, test, and keep-out validation, instead of evaluation patterns only on the same data used to construct them.

This distinction is important because station-level and connector-level logs do not always carry the same operational meaning. Some issues are not asset-network or station-contextual, while others are connector-specific and closer to the user-facing charging experience. Therefore, the methodology avoids collapsing all signals into one generic charger-level view. This is particularly important for DC chargers, where some issue signals are more visible at connector-level and may be hidden if only station-level logs are used.

The data also has important limitations. It does not consistently contain confirmed root-cause labels, technician findings, or repair outcomes. Therefore, issue-family flags are treated as operational proxies rather than verified failure labels. Telemetry coverage is also uneven, especially for voltage and current data, so missingness and telemetry availability are treated as part of the data reality rather than ignored.

Charging-session data is useful for user-impact interpretation, but a short or failed session may be caused by charger, user, authorization, communication, or vehicle-side factors. EV-WATTS similarly uses charging-session data, error codes, and power-transfer data to evaluate charging reliability, supporting the need to interpret status logs together with session outcomes rather than in isolation (CLEAResult, 2024). **Table 1** shows data layers that are used in the analysis. Further details on dataset structure, station-day and connector-day construction, and cleaning rules are provided in **Appendix 6.1 – 6.2**.

Table 1. Data layers used in the analysis.

Data layer	Main role in the analysis	Key limitation
Station status logs	Station-level availability and error signals	May hide connector-specific issues
Connector status logs	Connector/session-facing error signals	More granular and noisy
CRD/session records	Session quality and user-impact proxy	Failed sessions can have multiple causes
Meter Values/telemetry	Voltage/current context where available	Uneven coverage and partly missing
Charger metadata	Vendor, protocol, firmware, age, asset context	Contextual, not causal by itself
Vendor diagnostic fields	Vendor Error Code and Info interpretation	Not standardized across vendors

2.5.2 Methodology

The methodology was designed as a staged analytical pipeline rather than a direct jump into predictive modelling. This is because raw charger data is noisy, repeated, vendor-dependent and not always diagnostic. In addition, the target variable is not a verified repair label, but an issue-family proxy derived from operational status and diagnostic signals.

The selected methodology therefore combines rule-based screening, issue-family feature engineering, and baseline machine-learning validation. **Table 2** highlights the selected analytical pipeline. The broader method exploration and selection logic behind this pipeline are provided in **Appendix 6.3**.

Rule-based thresholds were used as screening heuristics, not as final operational alert rules. They helped distinguish high-volume signals from signals that were more concentrated, persistent, or fault linked. This was necessary because high frequency does not automatically mean high severity. Some error codes behave more like warning-like or contextual signals, while others are strongly linked to Faulted states. ChargeX's analysis on minimum required error codes also highlights that inconsistent error reporting complicates diagnosis and reliability improvement across charging systems, supporting the need to structure raw error messages into more interpretable families before using them for maintenance decisions (Savargaonkar et al., 2023). Detailed screening rules, issue-family mapping, feature-role classification, and baseline modelling setup are provided in **Appendix 6.4 – 6.7**.

Table 2. The selected analytical pipeline methodology.

Stage	Purpose	Output
1. Data audit and panel construction	Align events, sessions, telemetry, and metadata records	Station-day and connector-day panels
2. Raw signal and filtered issue analysis	Distinguish raw frequency from operational relevance	Error footprint and faulted-link checks
3. Issue-family mapping	Translate raw Error Code, Vendor Error Code, and Info into interpretable families	Issue-family flags
4. Feature engineering	Build recurrence, session-quality, exposure, telemetry, metadata, and diagnostic features	Candidate maintenance indicators
5. Feature screening and role classification	Separate early-warning candidates from same-day diagnostic/context features	Robust feature shortlist
6. Baseline issue-likelihood modelling	Test whether features support ranking outside the training period	Explainable ranking evidence

Issue-family mapping was then used to preserve technical meaning while reducing raw log complexity. Instead of treating each row error label separately related to signals were grouped into families such as metering/P1/Modbus, grounding/electrical protection, connector lock, voltage/grid quality, stack/power-module availability, EV

communication, authorization/session-start, emergency-stop/stop-loop. For broad categories such as Other Error, the mapping relied on additional vendor-specific diagnostic information from Vendor Error Code and Info.

Feature engineering was treated as a core methodological step because raw OCPP/status logs, charging session records, telemetry, and metadata are not directly model-ready. They first need to be transformed into operationally meaningful predictors, such as recurrence, session-quality indicators, exposure controls, telemetry context, metadata context, and diagnostic features. This follows feature-engineering literature, which emphasizes that predictive performance depends strongly on how raw data is represented as useful model inputs, not only on the selected algorithm (Zheng & Casari, 2018; Kuhn & Johnson, 2019). Same-day fault/status variables were treated carefully: they are useful for dashboards and troubleshooting explanation, but may cause leakage if used as early-warning predictors without temporal shifting. Therefore, feature screening classified features by role: early-warning candidate, diagnostic context, exposure control, telemetry validation, or dashboard-only support. Also, feature screening combined descriptive comparison, standardized mean difference, and Spearman rank correlation. Spearman correlation was used as a non-parametric association measure because many operational features are skewed, count-based, or rank-like. The screening results were used to classify features by role, not to claim causality.

Furthermore, baseline logistic regression was selected as the first modelling approach because it provides a transparent benchmark for tabular, issue-family classification under limited label maturity. It was not selected because it is the most advanced predictive-maintenance model, but because it tests whether the constructed features carry useful operational signal. This aligns with the broader argument that interpretable models should be preferred when they are sufficient for high-stakes or operational decision-support settings, rather than starting with black-box models and explaining them afterwards (Rudin, 2019).

The model output is interpreted as issue-likelihood ranking, not as calibrated fault probability. This distinction matters because issue-family days are relatively rare and the dataset is imbalanced. For imbalanced binary classification, precision-recall metrics are often more informative than Receiver Operating Characteristic (ROC)-based interpretation because they focus on the positive class, which is the class of operational interest (Saito & Rehmsmeier, 2015). In this project, PR-AUC, lift, test/keep-out stability, and operational interpretability are more relevant than accuracy alone. The purpose is to rank which station-day or connector-day should be reviewed first, for which issue family, and why.

2.6 Analytical Findings

This section presents the main analytical findings from the GreenFlux platform data. The purpose is not to report every technical output, but to identify which patterns matter for GreenFlux's maintenance intelligence direction. The analysis shows that charger-health interpretation cannot rely on raw error-code counts alone because frequency, severity, entity level, and diagnostic meaning do not always align. A high-volume signal may behave like a warning or context signal, while a lower-volume signal may be more strongly linked to Faulted status or recurring maintenance relevance.

The findings are structured in five steps. First, AC and DC chargers are compared to show why their operational behavior and predictive-maintenance logic should be separated. Second, station and connector-level status logs are examined to explain why raw OCPP error codes need further interpretation before they can support maintenance decisions. Third and fourth, Alfen AC and Alpitronic DC are used as representative case analyses to identify which issue families are most suitable for explainable issue-likelihood ranking. Finally, the section summarizes what these findings mean for GreenFlux's product direction.

The section should therefore be read as an analytical synthesis rather than a technical output list. Detailed status-code tables, hardware-oriented error explanation, raw Error Code entity-level hypothesis testing, and full AC/DC analytical results are provided in **Appendix 7**.

2.6.1 Different Behavior Between AC and DC, and Why Predictive Maintenance is Different

The first finding is that AC and DC charging should not be treated as one homogeneous operational category. The session data already shows different user and operational behavior. AC sessions dominate the dataset and show a wider duration pattern, including longer sessions that are consistent with destination, workplace, residential, or overnight-style charging. Compared with AC, the observed DC session in this dataset is fewer and more concentrated in shorter charging durations, reflecting a different user need and charger operating context, as can be seen in **figure 2**.

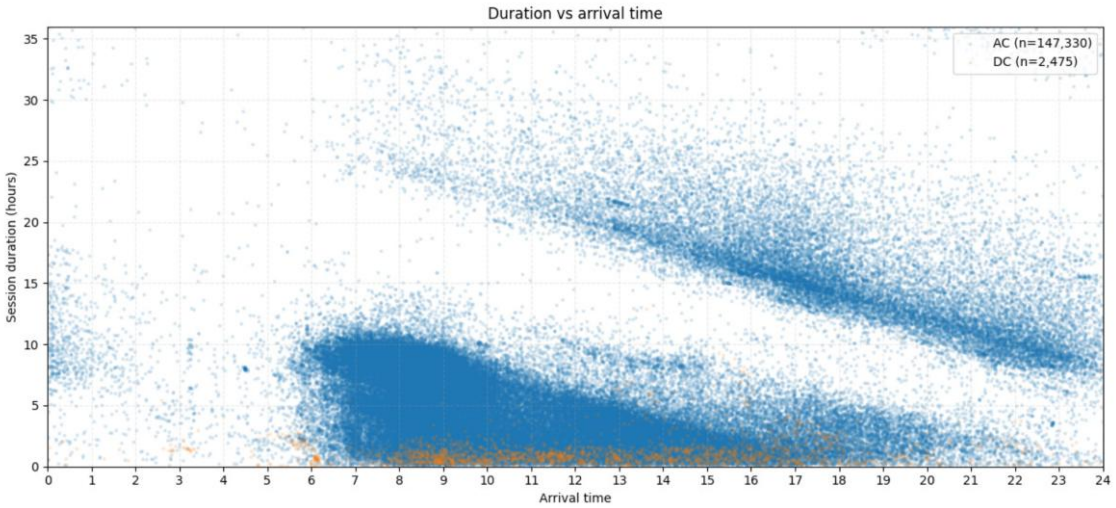

Figure 2. AC and DC sessions differ in duration and arrival-time patterns.

This figure supports the operational rationale for separating AC and DC before interpreting session-quality indicators and maintenance signals. For example, a short AC session may be more suspicious if it deviates from typical longer-duration behavior, while a short DC session can be closer to normal fast-charging behavior depending on context. This means that session-derived features such as short-session rate, valid-session share, and energy-per-session need to be interpreted differently for AC and DC chargers.

This broader reliability perspective is consistent with field evidence showing that public charging non-functionality can result from multiple causes, including charge-initiation failures, network failures, payment problems, screen issues, and broken

connectors, not only internal charger faults (Rempel et al., 2022). Public charging reliability studies and policy briefings similarly emphasize that charger reliability needs to be assessed from the user and operational experience, not only from backend availability status (Bernard, 2023).

The difference is not only behavioral. AC and DC chargers also differ in technical architecture, error semantics, and likely troubleshooting routes. AC chargers are more often linked to installation, metering, grid-quality, connector-lock, and electrical-protection contexts. DC chargers involve more complex power-module allocation, connector-level sequence control, vehicle-charger communication, cooling, metering/pairing, and charger-specific diagnostics. As a result, a single generic predictive-maintenance model would risk flattening important operational differences. **Table 3** stresses why AC and DC require different maintenance logic.

Table 3. Why AC and DC require different maintenance logic.

Dimension	AC (Alfen)	DC (Alpitronic)	Implication for maintenance intelligence
Session behavior	Higher observed session count, longer and more varied session duration	Fewer observed sessions, shorter and more concentrated session durations	Session-quality indicators should not be interpreted identically
Technical architecture	More installation, grid, metering, and connector-interface context	More power-module, sequence connector, and vehicle communication complexity	Issue-family logic should be charger type specific
Status/error semantics	Standard errors plus Other Error enrichment	Connector-led and Other Error-heavy	Error interpretation should be charger-type and vendor-specific
Operational response	Remote support, installation/electrical checks, or hardware engineer	Vendor guideline, connector-specific diagnosis, technician/site intervention	Troubleshooting support should suggest different action routes

These findings directly shape the rest of the analysis AC and DC should share a common maintenance-intelligence framework, but not the same interpretation logic. For GreenFlux, the practical implication is to develop charger type-specific issue families and feature sets rather than a single generic charger-health model.

2.6.2 Understanding Status Logs and Error Codes at Station and Connector Level

The second finding is that raw Error Code behavior already carries important entity-level information. Before mapping errors into issue families, the analysis first tested whether selected raw error codes appeared more meaningfully at station level, connector level, or both. This distinction matters because a station-level signal may

indicate asset/network-wide, site, cabinet, communication, or system-level context. A connector-level signal is often closer to user-facing charging behavior, such plug interaction, connector lock, session start, power transfer, or vehicle-charger communication. Some raw errors require both views because their operational meaning depends on whether the signal is isolated to one connector or reflected at the station level.

This station-versus-connector hypothesis testing supports the decision to preserve entity level in analytical design. Aggregating all signals into one generic charger-level view can hide actionable detail. Therefore, the analysis uses station-day and connector-day tables/panels rather than forcing all status and error events into one generic asset score. Only after this raw station-versus-connector behavior was assessed were signals mapped into issue families using Error Code, Vendor Error Code, Info, temporal recurrence, and session context. **Table 4** shows raw error code entity-level interpretation before issue-family mapping.

Table 4. Raw Error Code entity-level interpretation before issue-family mapping.

Raw Error Code Interpretation	Example signal	Operational meaning	Product implication
More station-contextual	Cabinet/door messages, station availability, some voltage/grid or metering context	Network-wide, site, cabinet, supply, communication, or system condition	Keep station-level health context
More connector-facing	Connector Lock Failure, selected connector-interface or session start signals	Plug/socket, connector, local current path, session-facing issue	Build connector-level ranking
Both/cross-level	Grounding, voltage/grid, metering depending on charger type and reporting behavior	May reflect either site/station condition or connector-specific impact	Preserve station-connector comparison
Under-specified raw code	Other Error or broad generic label	Raw error label does not explain the operational meaning	Use Vendor Error Code, Info, and issue-family mapping.

OCA's uptime monitoring guidance also supports this entity-aware interpretation. In OCPP 1.6, uptime monitoring relies mainly on Status Notification, status, and error code, while OCPP 2.0.1 introduces richer events and component-level monitoring through *Notify Event* and connector-status logic (Open Charge Alliance [OCA], 2024). This matters because the current data structure requires careful interpretation of status and error fields before they can be used as maintenance indicators.

A related finding is that Error Code should not be interpreted as a binary fault label. Some error codes behave like warning-like or contextual signals because they appear in

non-Faulted states and are weakly linked to subsequent Faulted status. Others are more fault-linked because they frequently appear together with Faulted or are shortly followed by a Faulted state.

For this reason, the analysis uses same-row Faulted overlap and Error-to-Faulted windows to distinguish warning-like/context signals from stronger triage candidates. This is important because high-frequency warnings may be less operationally urgent than lower-frequency fault-linked signals. **Table 5** highlights the interpretation on operational severity of raw error signals.

Table 5. Operational severity interpretation of raw error signals.

Error behavior	Evidence logic	Operational meaning	Product treatment
Warning-like/ context signal	Low same-row Faulted overlap and weak Error-to-Faulted link	Signal may indicate context, warning, retry, or transitional state	Monitor through dashboard or recurrence
Fault-linked signal	High same-row Faulted overlap or Error-to-Faulted link	More operationally severe and closer to maintenance relevance	Use as ranking and triage candidate
Diagnostic container	Broad label requires Vendor Error Code or Info interpretation	Raw label hides multiple possible meaning	Unpack before ranking or alerting

The analysis also shows that raw Error Code alone is often insufficient. Broad labels such as Other Error can contain very different operational meanings depending on Vendor Error Code, Info, status sequence, and charger type. ChargeX's minimum required error-code literature supports this problem at industry level: inconsistent error reporting makes charging diagnostics and resolution harder across charging systems, which creates a need for more structured error interpretation (Savargaonkar et al., 2024). For this reason, Other Error is treated as a diagnostic container, not as residual noise.

Furthermore, **figure 3** illustrates the operational interpretation around one of the error codes received from OCPP so-called Connector Lock Failure, and it shows why raw error-code interpretation needs additional diagnostic context. At first glance, it appears to be a hardware-oriented connector error. However, Alfen documentation and the dataset suggest that it can represent different operational situations. Vendor Error Code 404 is more consistent with a likely user or cable-attempt issue, where the cable could not be locked and the user may retry (Alfen, 2024a). While Vendor Error Code 211 is more hardware-like because it refers to a lock motor that could not move (Alfen, 2024b).

Temporal behavior adds further context: 404 shows more short-burst and same-day retry behavior, while recurrent 211 events are stronger maintenance candidates. However, this split should be treated as an indicative classification, not a confirmed root cause. It needs field-service records or technician findings to confirm the actual cause. This example supports a broader conclusion. Hardware-oriented error analysis should not rely on raw error labels alone. It should combine entity level, vendor-specific diagnostics, temporal recurrence, and operational outcomes.

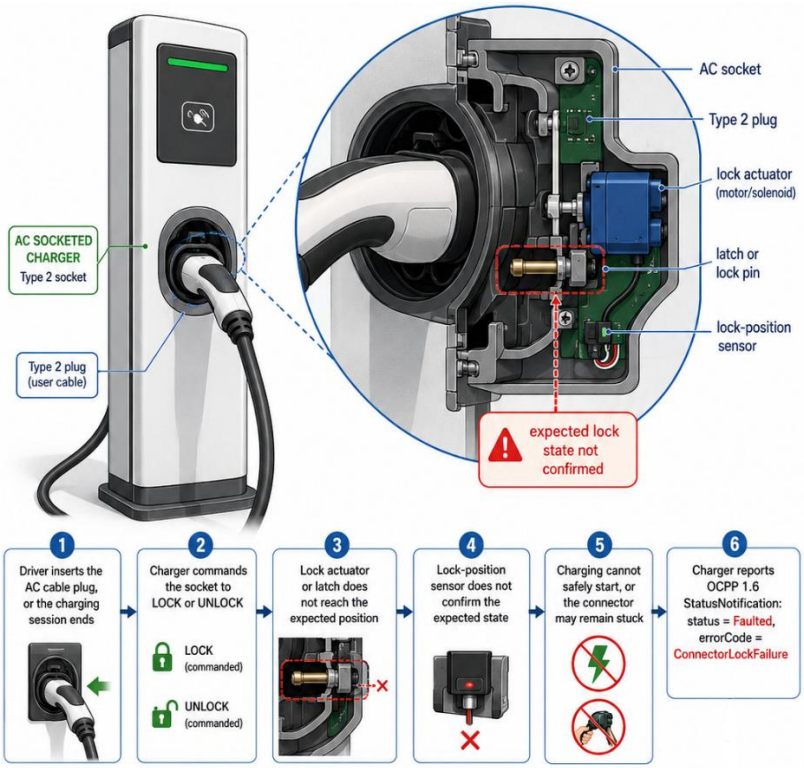

Figure 3. Operational interpretation around Connector Lock Failure error code in Alfen AC charger.

For GreenFlux, this creates a direct product implication. Maintenance intelligence should avoid flat charger-health score that aggregates all signals too early. A more useful layer should preserve station/connector context, map raw errors into issue families, distinguish warning-like from fault-linked signals, and explain why a station or connector is prioritized. Further details on raw Error Code entity-level hypothesis testing, warning-like versus fault-linked signal behavior and hardware-oriented error mechanisms are provided in **Appendix 7.1** and **Appendix 7.2**.

2.6.3 Alfen AC Representative Analysis

The Alfen analysis represents the AC charger case in this project. It applies the station/connector distinction and issue-family logic discussed earlier to identify which AC issue families are most suitable for maintenance-intelligence development. The main finding is that not all AC hardware-oriented signals are equally useful for predictive maintenance. The practical distinction is between issue families that are ready for review prioritization and issue families that should remain as dashboard or validation context.

For Alfen AC representative, the strongest issue-family candidates are **metering/ P1/Modbus/measurement** and **grounding/electrical protection**. Metering-related signals are operationally important because they can affect measurement reliability, session records, CDR quality, billing, and settlement. Grounding/electrical protection is closer to safety and protection logic, where recurring signals may require electrical or field inspection.

The model results should be interpreted as evidence for **explainable review prioritization**, not confirmed physical root-cause diagnosis. Because issue-family days are rare, PR-AUC assesses whether issue-days are concentrated among high-ranked

cases, while lift shows whether top-ranked assets contain more issue-days than random selection. Test and keep-out results indicate whether this ranking signal remains stable outside the training period. **Table 6** summarizes the Alfen AC evidence and operational use.

Table 6. Alfen AC issue-family evidence and operational use.

AC issue family	Entity level	Evidence interpretation	Recommended use
Metering/P1/Modbus/ Measurement	Station	Strongest AC signal. Ranking remains stable in test and keep-out with high PR-AUC and strong lift	First candidate for ranking and dashboard
Grounding/ electrical protection	Connector	Strong connector-level signal. Useful for prioritizing safety/protection-related review	Ranking and operational triage
Metering/P1/Modbus/ Measurement	Connector	Promising signal. Stable enough for ranking support, but weaker than station-level metering	Ranking and validation
Voltage/grid quality	Connector	Moderate signal. Operationally relevant, but should be validated with voltage/current telemetry	Telemetry validation and prioritization
Switching/component	Connector	Moderate ranking value, but needs stronger validation and field feedback	Ranking support
Connector lock/ mechanical interface	Connector	Useful for review when recurrent, but not strong enough for automatic alerting	Prioritization only
Overcurrent/ load protection	Connector	Useful for review when recurrent, but not strong enough to automatic alerting	Prioritization only
Station-level voltage/grounding	Station	Unstable or limited validation evidence	Context only/interpret cautiously

In this table, “ranking” means that the issue-family is suitable for prioritizing which station or connector should be reviewed first. “Dashboard” means that the signal is useful for operational visibility, while “validation” means that the signal should be combined with telemetry, ticket, or field evidence before being used for automated alerting.

The strongest numeric evidence comes from station-level metering/P1/Modbus/measurement, with PR-AUC of 0.887 in the test period and 0.916 in the keep-out period. Connector-level grounding/electrical protection also shows stable ranking value, with PR-AUC of 0.647 in the test period and 0.668 in the keep-out period. Connector-level metering/P1/Modbus/measurement is promising, with PR-AUC of 0.5333 in the test period and 0.611 in the keep-out period. These values should be

interpreted against the low baseline prevalence of issue-family days. Under class imbalance, the insight is not whether the model is accurate in a general sense, but whether it ranks relevant issue-days higher than random selection and remains stable across test and keep-out periods. Full model metrics, including ROC-AUC, lift, precision/recall, recurrence benchmarks and feature-screening results are provided in **Appendix 7.3**.

The AC results also show that recurrence is a dominant early-warning signal. Previous-day and rolling 7-day issue history repeatedly contribute to issue-likelihood estimation. Operationally, this is useful because recurring issue-family patterns can help prioritize assets for review. However, it also means the baseline model should not be overclaimed as discovering a complex causal mechanism. In many cases, the model is mainly learning that unresolved or recurring issue patterns are likely to reappear.

The distinction is important for product interpretation. Same-day diagnostic or Faulted signals are useful for dashboards and troubleshooting explanation, but they should not automatically be treated as early-warning predictors unless they are shifted into a prior window. For this reason, the AC model results are best interpreted as explainable issue likelihood ranking, not as calibrated fault prediction. The output can help GreenFlux and CPOs decide which station or connector should be reviewed first and why, but it does not yet prove the physical root-cause of the issue.

The AC case therefore suggests a staged maintenance-intelligence approach. Metering/P1/Modbus and grounding/electrical protection are the strongest first candidate for issue-family dashboarding and ranking. Voltage/grid quality and switching/component signals are operationally useful, but they need additional validation, especially through telemetry and field feedback. Connector lock and overcurrent remain relevant for prioritization, but their current model evidence is not strong enough for automation alerting.

Overall, Alfen AC shows that predictive maintenance development should begin with interpretable, issue-family specific ranking rather than a generic AC charger-health score. The strongest near-term value lies in identifying recurring metering and grounding patterns, linking them to session-quality and entity-level context and presenting them as explainable review priorities.

2.6.4 Alpitronic DC Representative Analysis

The Alpitronic analysis represents the DC charger case in this project. Compared with the Alfen AC case, the DC analysis shows a more connector-led and charger-specific diagnostic pattern. Many meaningful DC signals are not captured well by standard OCPP Error Code alone. Instead, they need to be reconstructed through connector-level logs and additional diagnostic information from Other Error, Vendor Error Code, and Info.

The main finding is that stack/power-module availability is the strongest DC candidate for explainable issue-likelihood ranking. Operationally, this issue family indicates that a connector may not have sufficient available internal power-stack capacity to support the expected charging process. This is specific to DC fast charging because DC chargers rely on internal power modules or stack to allocate available power to connectors and active charging sessions. When diagnostics such as No Free Stacks,

No Free Power, or related power-module availability messages appear, they can indicate a power-allocation or module-availability constraint rather than a generic charger fault.

As in the AC case, the model results should be interpreted as operational ranking evidence, not as calibrated failure probability. The purpose is to assess whether engineered features can help prioritize connector-days for review. In the DC case, the strongest result comes from stack/power-module availability, which shows PR-AUC of 0.623 in the test period and 0.589 in the keep-out period, with lift of approximately 7.37x in the test period and 6.80x in the keep-out period. In practical terms, this means that the model is able to concentrate stack/power-module issue-days among the top-ranked connector-days substantially better than random selection. **Table 7** highlights Alpitronic DC issue-family evidence and operational use.

Table 7. Alpitronic DC issue-family evidence and operational use.

DC issue family	Entity level	Evidence interpretation	Recommended use
Stack/power module availability	Connector	Strongest DC signal. Stable test and keep-out ranking evidence with the highest PR-AUC among DC candidates and strong lift.	First DC candidate for ranking and monitoring
EV communication/ session	Connector	Moderate signal. Relevant for session-start and authentication-related diagnostics.	Ranking and operational triage
Emergency stop/ stop-loop	Connector	Moderate signal. Useful for abnormal stop-loop, emergency-stop, or premature-unplug review.	Operational prioritization
Grounding/ electrical protection	Connector	Moderate signal with high lift, but relatively low positive count.	Ranking and validation
Metering/ pairing diagnostic	Connector	Operationally important and fault-linked, but weaker as a baseline ranking family.	Investigate and prioritize
Other diagnostic families	Connector	Includes safety/insulation/relay, EV voltage/current compatibility, sequence/controller state and unresolved vendor-specific diagnostics. Current evidence is exploratory or requires deeper feature engineering.	Manual review/ future feature engineering

The second group of DC candidate consists of EV communication/session-context, authorization/session-start, emergency-stop/stop-loop, and grounding /electrical protection. These candidates show moderate baseline ranking evidence. They are useful for prioritizing connector-days for diagnostic review, but they should not be converted into automatic alerting rules without stronger validation. They may reflect a mix of charger behavior, vehicle interaction, user behavior, communication flow, backend authorization, or abnormal session termination.

The more nuanced result is metering/pairing diagnostic. Earlier diagnostic analysis shows that this issue is operationally meaningful and fault linked. However, the baseline ranking model does not yet show prioritization evidence as strong as stack/power-module availability, authorization, or EV communication. This does not mean that metering/pairing is unimportant. Rather, it suggests that the current feature set is not specific enough. This is an example where diagnostic severity and model readiness do not fully align. A signal can be operationally important but still require more specific features before it becomes a strong ranking candidate. Future analysis may need richer Meter Values alignment, CDR meter-delta checks, and connector-meter consistency features.

The DC findings also confirm that recurrence remains a dominant early-warning signal. Previous-day and rolling-7-day issue history repeatedly contribute to issue-likelihood ranking. Operationally, this is useful because recurring connector-level diagnostic patterns can help prioritize review. However, as with the AC case, recurrence should not be over-interpreted as causal prediction. It may reflect unresolved diagnostic persistence rather than a newly emerging failure mechanism.

The key DC-specific implication is that Alpitronic maintenance intelligence should be connector-led and charger-specific. Here, charger-specific diagnostic interpretation means using Alpitronic-specific Vendor Error Code and Info messages to interpret broad OCPP categories such as Other Error. Unlike AC, where some hardware-oriented signals can appear meaningfully at both station and connector level, the DC analysis shows that many important diagnostic candidates are reconstructed at connector level, especially through Other Error + Vendor Error Code + Info. For GreenFlux, this means that a DC maintenance layer should not rely on standard OCPP Error Code or station-level summaries. It should preserve connector-level context, unpack Other Error diagnostics, and use issue-family-specific ranking to prioritize review.

Overall, Alpitronic DC shows that predictive-maintenance development should begin with connector-level issue-family ranking rather than a generic DC charger-health score. The strongest first candidate is stack/power-module availability, followed by moderate ranking candidates such as EV communication, authorization/session-start, emergency-stop/stop-loop, and grounding. Metering/pairing should remain in the timeline as a strategic diagnostic candidate, but it needs deeper issue-specific, feature engineering before it can be treated as a strong ranking family. Full DC issue-family mapping, Other Error unpacking, feature-screening outputs, recurrence benchmarking, and model-performance tables are provided in **Appendix 7.4**.

2.6.5 Summary of Analytical Findings

The analytical findings show that GreenFlux can move beyond reactive log inspection, but the current evidence supports a staged approach. The strongest near-

term opportunity is not fully automated predictive maintenance, but explainable issue-likelihood ranking, identifying which station or connector should be reviewed first, what issue-family is likely involved, and why the case deserves attention. **Table 8** highlights summary of analytical findings and strategic meaning.

Table 8. Summary of analytical findings and strategic meaning.

Finding	Evidence from analysis	Strategic meaning for GreenFlux
AC and DC behave differently	Session behavior, charger architecture, status/error semantics, and troubleshooting routes differ	Use one maintenance-intelligence framework, but separate AC/DC interpretation logic
Station and connector logs are not interchangeable	Raw Error Code testing shows station-contextual, connector-facing, and cross-level signals	Preserve station-day and connector-day logic before aggregation
Raw Error Code is not enough	Generic codes can hide different meanings, warning-like and fault-linked signals behave differently	Avoid raw error-count alerting; classify signals by entity level, severity behavior, and issue family
Other Error can be diagnostic	AC and DC findings show useful information in Vendor Error Code and Info	Treat Other Error as a diagnostic container, not residual noise
First ranking candidates differ by charger type	**AC:** metering/P1/Modbus and grounding **DC:** stack/power-module availability	Start ranking with issue families that have stronger evidence
Current evidence supports ranking, not calibrated prediction	Models rank issue-days better than random selection, but lack confirmed repair labels	Use explainable issue-likelihood ranking before automated predictive alerts

Overall, these findings suggest that GreenFlux should avoid a flat, generic charger-health score that aggregates all signals too early. A more defensible direction is an explainable maintenance-intelligence that preserves charger type, station/connector context, issue-family meaning, recurrence, session impact, and diagnostic evidence. This defines the boundary for the next section: GreenFlux can operationalize ranking and troubleshooting support in the near term, while calibrated predictive alerts require stronger feedback loops, field validation, and repair-outcome evidence.

2.7 Strategic Synthesis and Discussion for GreenFlux

This section integrates the external-environment, internal-position, competitor, operational-context, and analytical findings into a strategic interpretation for GreenFlux. The purpose is to discuss what the combined evidence means for GreenFlux's position in the CPMS market and for the development of charger reliability features. The central synthesis is that GreenFlux has a credible opportunity to move from reactive log monitoring toward charger health and uptime intelligence. However, the evidence also shows that this opportunity should not be framed too early as fully automated predictive

maintenance. A more defensible strategic direction is explainable uptime intelligence: a capability that helps CPO customers interpret operational signals, prioritize station or connector review, and understand the evidence behind prioritization.

2.7.1 Synthesis of External, Internal, and Analytical Findings

The external analysis shows that charger reliability is becoming more visible, measurable, and strategically important. Regulation increases accountability for uptime and service quality, CPO economics increases the need for operational efficiency, and EV users increasingly expect public charging to be dependable. Stakeholder analysis further shows that charger reliability is not controlled by one actor. CPOs, software platforms, charger OEMs, field-service partners, site and grid-related actors, regulators, and EV users all shape whether charging services are reliable in practice. This creates pressure on CPMS providers to move beyond basic connectivity and monitoring toward operational support that helps CPOs respond faster and more consistently.

The internal analysis shows that GreenFlux has a strong platform position, but that position is not automatically a sustained advantage. GreenFlux operates close to relevant operational data flows, including OCPP status logs, connector status events, session records, vendor-specific information, and platform-side operational signals. Its experience in CPMS operations, interoperability, roaming, smart charging, and customer support gives it a credible basis for reliability-oriented feature development. However, the VRIO and 7S assessments suggest that the defensible advantage still needs to be built. Data access alone is not rare in the CPMS market, and maintenance intelligence is not yet established as a mature core product layer. The stronger internal opportunity lies in converting fragmented platform signals into workflow-compatible interpretation, prioritization, and support.

The competitor and best-practice analysis reinforce this point. CPMS differentiation is moving beyond basic monitoring, billing, roaming, and status visibility. Market claims around AI, self-healing, diagnostics, and predictive maintenance are becoming more common, but public evidence varies in specificity and validation strengths. This creates both competitive pressure and strategic space. GreenFlux does not need to compete by making broad AI claims. A stronger position would be to build a credible and explainable reliability layer that is clearer about what the system can detect, what it can prioritize, and what still requires human or field validation.

The operational and data analysis explains why this caution is necessary. OCPP status logs, error codes, vendor information, and charging-session records are useful, but they are not equivalent to confirmed failure diagnosis. Availability and charging success are not the same thing. A charger may appear available while still producing failed or abnormal user experience because of authorization, communication, connector, vehicle-side, payment, or power-transfer issues. The analysis also shows that charger type and entity level matter. AC and DC chargers do not produce the same operational patterns, and station-level and connector-level logs should not be aggregated too early. A generic charger-health score may be attractive for simplicity, but it risks hiding the diagnostic detail needed for maintenance decisions.

The analytical findings confirm this interpretation. For Alfen AC, metering/P1/Modbus/measurement and grounding/electrical protection emerge as strong first candidate for operational prioritization. For Alpitronic DC, stack/power-

module availability is the clearest first candidate, while authorization/session-start, EV communication, and other session-flow families provide more moderate prioritization value. These findings do not prove physical root causes. They show that selected issue families can support explainable review prioritization when interpreted with entity level, recurrence, session impact, vendor context, and validation evidence. The strategic interpretation is therefore not that GreenFlux is ready for autonomous predictive maintenance, but that it has enough evidence to develop a more structured maintenance-intelligence layer.

2.7.2 SWOT and TOWS: Strategic Options for GreenFlux

SWOT is used here to synthesize internal strengths and weaknesses with external opportunities and threats, while TOWS translates this synthesis into strategic options rather than final recommendations (Weihrich, 1982; Whittington et al., 2023). This synthesis clarifies the strategic tension facing GreenFlux: it has a valuable platform and data position, but the market is becoming more competitive, and the current evidence does not yet support broad predictive-maintenance claims. **Table 9** highlights the overview of SWOT analysis.

Table 9. The results of SWOT analysis.

SWOT element	Strategic synthesis
Strengths	GreenFlux has a strong CPMS position, access to charger status and session data, experience in multi-actor charging operations, interoperability capabilities, and a platform through which reliability features can be embedded into CPO workflows.
Weaknesses	Platform signals remain fragmented across vendors, charger types, protocol version, station/connector levels, and session contexts. Confirmed failure labels, support-ticket linkage, and field-maintenance feedback are not yet mature enough for fully automated prediction.
Opportunities	Charger reliability is becoming more important for regulation, CPO economics, user trust, and CPMS differentiation. There is room for a more credible reliability layer that translates raw charger data into clearer operational interpretation.
Threats	Generic CPMS competition, OEM diagnostic tools, AI diagnostics firms, uptime specialists, internal CPO analytical teams, alert fatigue, and overclaiming predictive capability can all weaken GreenFlux's differentiation if reliability features remain generic.

Furthermore, the TOWS synthesis, as can be seen in **table 10**, translates these findings into strategic options. The options do not yet prescribe a full implementation roadmap; rather, they clarify which strategic direction is most consistent with the evidence.

The SWOT-TOWS synthesis points to one dominant strategic direction: GreenFlux should not position the next step as generic AI-based predictive maintenance. The stronger direction is to build an explainable reliability layer that can mature toward predictive maintenance over time. This is more consistent with the current data maturity, competitor landscape, operational complexity, and stakeholder needs.

Table 10. The results of TOWS synthesis.

TOWS logic	Strategic option for GreenFlux
SO (Strengths-Opportunities): Use strengths to capture opportunities	Use GreenFlux's platform and data position to develop charger health and uptime intelligence as a reliability-oriented CPMS capability.
WO (Weaknesses-Opportunities): Address weaknesses to capture opportunities	Address fragmented diagnostics and weak feedback labels by starting with explainable review queues, issue-family interpretation, and structured operational feedback.
ST (Strengths-Threats): Use strengths to reduce threats	Differentiation against OEM tools, AI diagnostics firms, and generic CPMS rivals through vendor-aware, entity-aware, and workflow-compatible maintenance intelligence.
WT (Weaknesses-Threats): Reduce weaknesses and avoid threats	Avoid premature automation or black-box predictive claims. Begin with human-in-the-loop prioritization and confidence boundaries before moving toward calibrated predictive alerts.

2.7.3 Selected Strategic Direction: From Predictive Maintenance to Explainable Uptime Intelligence

The selected strategic direction is to move from predictive maintenance ambition toward explainable uptime intelligence as the near-term capability layer. Predictive maintenance can remain the longer-term ambition, but the current evidence supports a more staged and transparent approach. Fully autonomous predictive maintenance would require confirmed incident labels, repair outcomes, intervention feedback, calibration testing, and clearer links between early signals and verified failures. The current dataset is stronger for ranking and interpretation than for autonomous prediction.

In practice, explainable uptime intelligence should help operators answer four operational questions: what happened, where it happened, whether it matters, and what should be reviewed next. This means moving beyond basic monitoring without overclaiming automatic diagnosis. The system should classify likely issue families, preserve station- and connector-level interpretation, show evidence such as recurrence, session impact, faulted-link behavior, and vendor diagnostic context, and suggest an appropriate review direction such as monitoring, remote support, vendor escalation, or field inspection.

This framing avoids two weak extremes: basic monitoring leaves CPOs with too many raw logs, while fully automated prediction would overstate the current evidence. GreenFlux should therefore avoid one generic charger-health score and instead develop an entity-aware, issue-family-aware prioritization layer that preserves charger type, station/connector level, recurrence, session impact, vendor diagnostic evidence, and confidence boundaries. This makes explainable uptime intelligence the near-term foundation for future health indicators, calibrated alerts, and more automated predictive-maintenance functions once validation evidence and operational feedback become stronger.

Part III: Conclusion and Advice

3.1 Conclusion

The project investigated how GreenFlux can better enable CPO customers to maintain and troubleshoot charging stations more efficiently. The analysis shows that the main opportunity is not simply to collect more data or immediately build a black-box predictive-maintenance model. GreenFlux already has access to rich CPMS/OCPP data, but the operational value depends on how this data is interpreted, prioritized, and translated into troubleshooting support.

The answer to the main research question is therefore that GreenFlux should develop an explainable maintenance-intelligence. This layer should transform raw station, connector, session, vendor, and telemetry signals into issue-family interpretation and review prioritization. It should help operators answer four practical questions: what happened, where it happened, whether it matters, and what action route should follow. In this sense, the recommended direction is not immediate autonomous fault diagnosis, but structured decision support for charger health and uptime.

The findings support a staged product direction. AC and DC chargers require different interpretation logic; station-level and connector-level logs should not be aggregated too early; raw Error Code alone is insufficient without Vendor Error Code, Info, recurrence, entity level, and session context; and model outputs are currently more suitable for issue-likelihood ranking than calibrated predictive alerts. The strongest first candidates are AC metering/P1/Modbus/measurement, AC grounding/electrical protection, and DC stack/power-module availability. These findings provide a practical foundation for GreenFlux to move from reactive log monitoring toward explainable uptime intelligence.

3.2 Recommendations and Implementation Timeline

Building on the strategic synthesis in **Section 2.7**, the advice is to develop predictive maintenance as a staged product capability rather than as an immediate fully automated product. The first step should be a human-in-the-loop decision-support layer that makes charger-health interpretation more structured, transparent, and operationally useful. The sequence below should be read as an indicative implementation timeline, not as a fixed product commitment. The exact timing should follow GreenFlux's internal product cycle, data capacity, and CPO pilot availability. However, the sequence should remain staged: define the issue logic first, test explainable ranking next, validate workflow fit, build feedback loops, and only then move toward calibrated predictive alerts.

Recommendation 1: Build an issue-family dashboard/layer for selected AC/DC issue families

In the first **0-4 months**, GreenFlux should define and dashboard a narrow set of high-evidence issue families. The initial scope should include AC metering/P1/Modbus/

measurement, AC grounding/electrical protection, DC stack/power-module availability, and selected moderate DC session-flow issues such as EV communication, authorization/session-start, and emergency stop/stop-loop. In the first 0-2 months, the product team should lead the issue-family taxonomy and pilot scope, supported by data/engineering, support, and customer success. In months 2-4, data/engineering should translate raw Error Code, Vendor Error Code, Info, recurrence, entity level, and session context into issue families. The decision gate is whether product, data, and support teams agree on the issue-family definitions, confidence boundaries, and dashboard outputs for the first pilot scope.

Recommendation 2: Add station-level and connector-level review queues

In **months 4-6**, GreenFlux should add explainable station-level and connector-level review queue for the selected issue families. The review queue should not replace existing station and connector views. Instead, it should sit on top of them as a prioritization layer. The output should show which station or connector should be reviewed first, what issue family is likely involved, and which evidence supports the ranking. Product and data/engineering should co-own this stage, with support, customer success, and selected CPO users validating whether the ranking is understandable and useful. The decision gate is whether pilot users can understand why cases are ranked and use the queue for review prioritization without treating it as confirmed fault diagnosis.

Recommendation 3: Provide reason codes, model-quality labels, and troubleshooting routes

From **months 4-9**, GreenFlux should strengthen the review queue with reason codes, model-quality labels, and suggested troubleshooting routes. Each ranked case should include interpretable reason codes, such as recurrence, short-session patterns, lower valid-session share, faulted-link evidence, telemetry context, or vendor diagnostic messages. Model-quality labels should distinguish between strong, moderate, weak-but-useful, and dashboard-only issue families. From **months 6-9**, support and customer success should lead the validation of action routes, supported by product, selected CPO users, and vendor or field-service contacts. The decision gate is whether the suggested routes fit real support and CPO workflows and do not overstate diagnostic certainty.

Recommendation 4: Build a feedback loop before automated alerts

In **months 9-12**, GreenFlux should build the operational feedback loop needed for future predictive-maintenance maturity. Ranked cases should be linked to support tickets, remote reset/reboot results, CPO feedback, vendor diagnosis, technician findings, installation checks, false positives, and confirmed repairs. Product and data/engineering should co-own this stage, with support, customer success, CPOs, field-service partners, and vendors contributing outcome information. The decision gate is whether reviewed cases can be labelled with operational outcomes consistently enough to support model improvement, calibration testing, and false-positive analysis.

Recommendation 5: Establish cross-functional ownership for maintenance intelligence

From the start of the timeline, GreenFlux should treat maintenance intelligence not only as a data-science feature, but as a cross-functional product capability. GreenFlux should assign a clear internal product owner for maintenance intelligence,

supported by a working group across product, data/engineering, support, customer success, and CPO-facing teams. These teams need shared ownership over issue-family definitions, ranking outputs, reason-code wording, feedback capture, and escalation boundaries. Leadership and product should use the decision gates from each stage to decide when to expand from dashboard to review queues, from review queues to troubleshooting support, and eventually from troubleshooting support to calibrated predictive alerts.

Only after **12 months or more** should GreenFlux test whether selected ranked signals can be converted into calibrated probability-based alerts or escalation logic. This should proceed only if feedback labels, calibration performance, false-positive tolerance, and user trust are sufficient. If ticket linkage, field-service feedback, or vendor-confirmed outcomes are not immediately available, the first six months should remain focused on dashboarding and review prioritization rather than calibrated prediction. The first release should therefore be positioned as decision support, not autonomous maintenance automation. **Appendix 8** translates this staged direction into implementation-support material, including an MVP product concept, feedback-loop and ownership design, and expanded adoption-risk register.

3.3 Feasibility, Organizational Readiness, and Adoption Risks

The proposed direction is feasible as a staged development, but not as an immediate fully automated prediction product. The current data can support issue-family dashboards, recurrence-based indicators, entity-aware ranking, and reason-code explanation for selected issue families. However, calibrated predictive alerts require stronger data maturity, especially confirmed repair labels, ticket linkage, field feedback, and calibration testing. The key feasibility condition is therefore not only whether GreenFlux has enough data, but whether the organization can turn that data into a trusted and workflow-compatible capability.

Building on the 7S assessment in **Appendix 3** and the SWOT-TOWS synthesis in **Section 2.7**, the main organizational requirement is cross-functional alignment. Product teams need to define product scope and customer-facing outputs. Data and engineering teams need to build feature tables, ranking logic, and model monitoring. Support and customer success teams need to validate whether reason codes and review queues fit real troubleshooting workflows. CPO-facing teams and pilot customers need to test whether the outputs reduce issues. **Table 11** summarizes the ownership logic.

From a stakeholder perspective, the value proposition is operational clarity. CPOs benefit from faster triage and clearer escalation decisions, while GreenFlux benefits from a differentiated uptime-intelligence capability and a clearer pathway toward future machine-learning development. **Table 12** summarizes the expected stakeholder value and feasibility conditions.

The main adoption risks are false confidence, false positives, workflow mismatch, insufficient feedback, vendor-specific behavior, over-aggregation, and over-automation. These risks can be mitigated by positioning the first release as decision support rather than automatic maintenance automation, starting with high-evidence issue families, using model-quality labels, keeping operators in the loop, and collecting feedback before expanding into automated alerts. **Table 13** highlights the main risks and safeguards.

Table 11. Ownership model for maintenance-intelligence development.

Workstream	Primary responsibility	Key contributors	Main task
Product direction	Product team	Support, customer success, data/ engineering	Define pilot scope, issue-family taxonomy, user stories, and decision boundaries
Data and model pipeline	Data/ engineering team	Product, support	Build reusable feature tables, ranking logic, model-quality labels, monitoring, and validation checks
Troubleshooting workflow	Support/ customer success	Product, CPO-facing teams	Test whether reason codes, evidence summaries, and action routes fit real support and CPO workflows
Customer validation	CPO-facing teams	Selected CPO pilot users, support	Validate usefulness, clarity, false-positive tolerance, and adoption barriers
Vendor and field feedback	Support/ customer success	Vendors, field-service partners, CPO technical teams	Capture whether remote actions, vendor escalation, or field inspection resolve the issue
Governance and scaling	Product and leadership	Data/engineering, support, customer success	Decide when to expand from dashboarding to ranking, from ranking to playbooks, and from playbooks to calibrated alerts.

Table 12. Expected stakeholder value and feasibility.

Stakeholder	Expected value	Feasibility condition
CPO operations	Faster prioritization and less manual log inspection	Ranking must be explainable and linked to existing workflows
CPO technical teams	Better distinction between monitor, rest, escalation, and field inspection	Issue families need action-oriented definition
GreenFlux product team	Differentiated uptime-intelligence capability	First release should be narrow, evidence-based, and customer-testable
GreenFlux support/ customer success	More structured troubleshooting conversations	Outputs should cover reason codes, confidence boundaries, and raw evidence links
Data/engineering teams	Clearer Machine Learning development pathway	Feedback labels and calibration logic must be developed
Vendors and field-service partners	Clearer evidence trail for escalation and inspection	Issue-family logic should preserve vendor, charger type, and station/connector context

Table 13. Adoption risks and safeguards.

Adoption risk	Safeguard
False confidence and over-automation	Position the first release as decision support, not confirmed diagnosis. Use ranking scores, reason codes, model-quality labels, and confidence boundaries. Move to predictive alerts only after labels, calibration, and false-positive tolerance are validated
False positive and alert fatigue	Start with selected high-evidence issue families, keep human review before alert automation, and monitor whether operators find the queue useful rather than noisy
Workflow mismatch	Validate review queues and action routes with support teams and selected CPO users before scaling. Outputs should fit real troubleshooting and escalation process
Insufficient feedback labels	Link rankings with support tickets, remote actions, vendor diagnosis, field outcomes, false-positive reviews, and confirmed repairs
Vendor-specific behavior and over-aggregation	Preserve vendor, charger type, AC/DC, station/connector level, issue-family, recurrence, and session-impact context instead of using one generic health score

The practical implication is that adoption should be gradual. GreenFlux should first demonstrate the ranking improves review prioritization and troubleshooting clarity. Only after operational users trust the output should the capability move toward automated alerts or prescriptive maintenance recommendations.

3.4 Reflection and Limitations

This project provides a practical path toward maintenance intelligence, but the findings should be interpreted with clear boundaries. The analysis relies on issue-family indicators rather than confirmed physical root-cause labels. This means the results are suitable for prioritization and diagnostic support, but not for causal failure claims or automatic field-dispatch decisions.

Recurrence is also useful but should be interpreted carefully. Repeated signals may indicate unresolved issues rather than true early warning, so recurrence should be combined with session impact, entity level, vendor diagnostics, and operational feedback. The analysis also focuses on representative AC and DC cases: Alfen and Alpitronic. These cases reveal important differences between AC and DC behavior, but the findings should not be generalized to all charger vendors without further validation.

Finally, telemetry coverage and field feedback remain incomplete in the used datasets. Voltage, current, meter values, remote-action outcomes, technician reports, and support ticket would strengthen the analysis substantially. The next methodological step should therefore be to connect CPMS signals with operational resolution data. Overall, predictive maintenance should be treated as a staged capability and organizational data-maturity journey. For GreenFlux, the most realistic next step is explainable uptime intelligence: issue family dashboard, station/connector ranking, reason-code explanations, and feedback loops. This creates a defensible bridge between current CPMS monitoring and future predictive maintenance.

References

Amaliah, N. R., Tjahjono, B., & Palade, V. (2025). Human-in-the-loop XAI for predictive maintenance: A systematic review of interactive systems and their effectiveness in maintenance decision-making. *Electronics, 14*(17), 3384. https://doi.org/10.3390/electronics14173384

Adner, R. (2017). Ecosystem as structure: An actionable construct for strategy. *Journal of Management, 43*(1), 39–58. https://doi.org/10.1177/0149206316678451

Alfen. (2024a, June 4). *What does error code 211 mean?* Alfen Knowledge Base. https://knowledge.alfen.com/categories/CAT-00927/KA-01312

Alfen. (2024b, June 4). *What does errorcode 404 mean?* Alfen Knowledge Base. https://knowledge.alfen.com/categories/CAT-00927/KA-01322

Ampcontrol. (2024). *EV charger repair and predictive maintenance: Boosting reliability with predictive maintenance*. Ampcontrol.

AMPECO. (2026). *AMPECO issues detection: Proactive network monitoring from fault to resolution*. AMPECO.

Barney, J. (1991). Firm resources and sustained competitive advantage. *Journal of Management, 17*(1), 99–120. https://doi.org/10.1177/014920639101700108

Bernard, M. R. (2023). *Improving public charging infrastructure reliability*. International Council on Clean Transportation.

Bernard, M. R., Díaz, S., Kok, I., & Tietge, U. (2024). *Enhancing user experience in public electric vehicle charging*. International ZEV Alliance.

Chen, B., Moore, A., Jia, B., Zhang, K., & Cao, M. (2025). Equity considerations in public electric vehicle charging: A review. *World Electric Vehicle Journal, 16*(10), 553. https://doi.org/10.3390/wevj16100553

Carvalho, T. P., Soares, F. A. A. M. N., Vita, R., Francisco, R. P., Tavares Vieira Basto, J. P., & Gomes Soares Alcalá, S. (2019). *A systematic literature review of machine learning methods applied to predictive maintenance*. Computers & Industrial Engineering, 137, 106024.

Chargecloud. (2025). *What is a Charge Point Management System?* Chargecloud.

ChargePoint. (2025). *Unified EV charging management software*. ChargePoint.

Chirumalla, K., Parida, V., Sjödin, D., & Wincent, J. (2023). Moving from servitization to digital servitization: Identifying the microfoundations of digital servitization capabilities. *Industrial Marketing Management, 113*, 321–337. https://doi.org/10.1016/j.indmarman.2023.01.019

Çınar, Z. M., Nuhu, A. A., Zeeshan, Q., Korhan, O., Asmael, M., & Safaei, B. (2020). *Machine learning in predictive maintenance towards sustainable smart manufacturing in Industry 4.0*. Sustainability, 12(19), 8211.

CLEAResult. (2024). *Electric vehicle charging station reliability*. EV WATTS. https://www.clearesult.com/sites/default/files/2024-05/EV-WATTS_White-Paper_Charging-Station-Reliability.pdf

Cummins, L., Sommers, A., Bakhtiari Ramezani, S., Mittal, S., Jabour, J. E., Seale, M., & Rahimi, S. (2024). Explainable predictive maintenance: A survey of current methods, challenges and opportunities. *IEEE Access*.

D'Agostino, B., Chen, J., & Rajagopal, R. (2026). *Beyond uptime: Actionable performance metrics for EV charging site operators*. arXiv preprint arXiv:2601.10861.

Department for Transport. (2024, October 21). *Public Charge Point Regulations 2023 guidance*. GOV.UK.

Driivz. (2025). *From reactive repairs to predictive performance*. Driivz.

European Commission. (n.d.). *Energy performance of buildings directive*. Directorate-General for Energy.

European Commission. (n.d.). *Questions and answers on the regulation on the deployment of alternative fuels infrastructure (EU 2023/1804)*. Directorate-General for Mobility and Transport.

European Commission. (n.d.). *Alternative fuels infrastructure*. European Commission.

European Commission, Directorate-General for Competition. (2023). *Study on the electric vehicle recharging market across the EU27 and the UK*. European Commission.

EVailable.ai. (n.d.). *Boosting EV-charging availability*. EVailable.ai.

Federal Highway Administration. (2023). *National Electric Vehicle Infrastructure Standards and Requirements*.

GreenFlux. (n.d.). *Smarter EV charging software*.

Hsaini, S., Ghogho, M., & Charaf, M. E. H. (2022). An OCPP-based approach for electric vehicle charging management. *Energies, 15*(18), 6735. https://doi.org/10.3390/en15186735

International Energy Agency. (2026). *Global EV Outlook 2026*.

Jardine, A. K. S., Lin, D., & Banjevic, D. (2006). *A review on machinery diagnostics and prognostics implementing condition-based maintenance*. Mechanical Systems and Signal Processing, 20(7), 1483–1510.

Karmaker, A. K., et al. (2023). Multi-stakeholder perspectives for transport electrification: A review on placement and scheduling of electric vehicle charging infrastructure. *Journal of Cleaner Production, 427*, 139145.

Kuhn, M., & Johnson, K. (2019). *Feature engineering and selection: A practical approach for predictive models*. Chapman & Hall/CRC.

Li, S., et al. (2025). Deep learning predicts real-world electric vehicle direct current fast charging profiles and detects irregular charging behaviours. *Nature Communications*.

Martins, J. A., & Rodrigues, J. M. F. (2025). Intelligent monitoring systems for electric vehicle charging. *Applied Sciences, 15*(5), 2741. https://doi.org/10.3390/app15052741

Mendelow, A. L. (1981). Environmental scanning: The impact of the stakeholder concept. In *Proceedings of the International Conference on Information Systems* (pp. 407–418).

Monta. (2025). *Alerts*. Monta. https://monta.com/en/features/alerts/

National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)*. U.S. Department of Commerce. https://doi.org/10.6028/NIST.AI.100-1

Netherlands Enterprise Agency. (2023). *Dutch National Charging Infrastructure Agenda*. RVO.

Open Charge Alliance. (2024). *Improving uptime monitoring with OCPP* (Version 1.3). Open Charge Alliance. https://openchargealliance.org/ocpp-info-whitepapers/159297-2/

Open Charge Alliance. (2024). *Open Charge Point Protocol (OCPP)*. https://openchargealliance.org/protocols/open-charge-point-protocol/

Open Charge Alliance. (2024). *What is new in OCPP 2.0.1*. Open Charge Alliance.

Ozkaya, M., Turunç, A., & Togrul, Y. T. (2026). A reference architecture for smart charging management systems for electric vehicles. *Designs, 10*(1), 4. https://doi.org/10.3390/designs10010004

Rempel, D., Cullen, C., Matteson Bryan, M., & Vianna Cezar, G. (2024). Reliability of open public electric vehicle direct current fast chargers. *Human Factors, 66*(11), 2528–2538. https://doi.org/10.1177/00187208231215242

Porter, M. E. (2008). The five competitive forces that shape strategy. *Harvard Business Review, 86*(1), 78–93.

Road. (n.d.). *Modular platform for EV charging*. Road.

Roeckle, F., Amorim, M., Keicher, L., Kern, M., & Rupprecht, J. (2022). *"Charging process chain": Holistic approach for identifying weaknesses in the electric vehicle charging ecosystem*. In *Proceedings of the 35th International Electric Vehicle Symposium and Exhibition (EVS35)*, Oslo, Norway.

Rudin, C. (2019). *Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead*. Nature Machine Intelligence, 1, 206–215.

Saito, T., & Rehmsmeier, M. (2015). *The precision-recall plot is more informative than the ROC plot when evaluating binary classifiers on imbalanced datasets*. PLOS ONE, 10(3), e0118432.

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.

Savargaonkar, M., et al. (2023). *Recommendations for minimum required error codes for electric vehicle charging infrastructure*. ChargeX Consortium / Idaho National Laboratory.

Sharma, J., Mittal, M. L., & Soni, G. (2022). *Condition-based maintenance using machine learning and role of interpretability: A review*. International Journal of System Assurance Engineering and Management.

Singh, R., Quinn, C., & MacKenzie, D. (2025). Poor reliability of public charging stations can impede the growth of the electric vehicle market. *Transport Policy, 171*, 695–705. https://doi.org/10.1016/j.tranpol.2025.06.026

Stiglic, G., Kocbek, P., Fijacko, N., Zitnik, M., Verbert, K., & Cilar, L. (2020). *Interpretability of machine learning-based prediction models in healthcare*. Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, 10(5), e1379.

Teece, D. J. (2007). Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance. *Strategic Management Journal, 28*(13), 1319–1350. https://doi.org/10.1002/smj.640

van der Koogh, M., Chappin, E., Heller, R., & Lukszo, Z. (2023). Stakeholder prioritizations for electric vehicle charging across time periods. *Transport Policy, 142*, 173–189. https://doi.org/10.1016/j.tranpol.2023.09.003

Virta. (2024). *Self-healing: Boost reliability with automatic fault recovery for your EV chargers*. Virta.

Waterman, R. H., Jr., Peters, T. J., & Phillips, J. R. (1980). Structure is not organization. *Business Horizons, 23*(3), 14–26. https://doi.org/10.1016/0007-6813(80)90027-0

Webb, A., et al. (2023). Reliability of electric vehicle charging infrastructure: A cross-lingual deep learning approach. *Communications in Transportation Research*.

Weihrich, H. (1982). The TOWS matrix—A tool for situational analysis. *Long Range Planning, 15*(2), 54–66. https://doi.org/10.1016/0024-6301(82)90120-0

Whittington, R., Regnér, P., Angwin, D., Johnson, G., & Scholes, K. (2023). *Exploring strategy: Text and cases* (13th ed.). Pearson.

Wolbertus, R., Jansen, S., & Kroesen, M. (2020). Stakeholders' perspectives on future electric vehicle charging infrastructure developments. *Futures, 123*, Article 102610. https://doi.org/10.1016/j.futures.2020.102610

Wu, L., Liu, M., Gong, K., Jiao, L., Huo, X., Zhang, Y., & Wang, H. (2025). The usage of big data in electric vehicle charging: A comprehensive review. *Energies, 18*(19), 5066. https://doi.org/10.3390/en18195066

Yu, Q., Que, T., Cushing, L. J., Pierce, G., Shen, K., Kejriwal, M., Yao, Y., & Zhu, Y. (2025). Equity and reliability of public electric vehicle charging stations in the United States. *Nature Communications, 16*, 5291.

Zheng, A., & Casari, A. (2018). *Feature engineering for machine learning: Principles and techniques for data scientists*. O'Reilly Media.

Appendix

Appendix 1. PESTEL Analysis: External Pressure on Charger Health and Uptime Intelligence

This appendix provides the evidence based behind **section 2.1**. The PESTEL analysis is used to identify external trends and pressure that are directly relevant to GreenFlux's project goal, by enabling CPO customers to maintain and troubleshoot charge stations more efficiently (Whittington et al., 2023). The analysis focuses on the external developments that affect charger uptime, maintenance intelligence, and CPMS differentiation.

Political

EV charging policy is increasingly moving from deployment volume toward infrastructure quality, interoperability, user information, and payment accessibility. In the European Union, the Alternative Fuels Infrastructure Regulation has applied since 13 April 2024 and aims to ensure minimum infrastructure, full interoperability, comprehensive user information, and adequate payment options for alternative-fuel infrastructure (European Commission, 2024). This matters for CPOs because public charging is increasingly treated as a service infrastructure that must be usable, transparent, and accessible, not only physically deployed.

For GreenFlux, this creates a role in supporting CPOs with the software capabilities needed to manage charging infrastructure at scale. Reliable availability information, structured monitoring, and service-performance reporting are becoming part of the broader infrastructure-quality agenda. However, GreenFlux should not position itself as the sole actor responsible for charging reliability. Its strategic role is more specific on helping CPOs make charging operations more visible, measurable, and manageable through better use of platform data.

Economic

The economic pressure comes from the scaling of EV and charging networks. Global electric car sales exceeded 20 million in 2025 and reached around 25% of total car sales, while public charging infrastructure continued to expand in parallel (International Energy Agency [IEA], 2026). The global stock of public charging points reached more than 7 million by the end of 2025, with China alone accounting for more than 65% of public charging points globally (IEA, 2026). As network scale, downtime becomes more economically relevant because each unavailable charger can reduce utilization, create missed charging opportunities, and increase operational workload.

For CPOs, reliability is therefore not only a technical matter but also an asset-performance issue. A charger that is installed but frequently unavailable creates limited commercial value and may increase support costs. For GreenFlux, this creates a clear

business implication. Uptime-related features should not be framed only as technical troubleshooting tools. They can become business-performance features that help CPOs reduce avoidable downtime, repeated manual log inspection, inefficient triage, and unnecessary field intervention. The business value, however, depends on whether these features fit CPO workflows and lead to decisions that operators can act on.

Social

Charging reliability affects whether EV adoption feels practical for users. Public charging user-experience research identifies availability, reliability, and interoperability as key dimensions of charging usability (Bernard et al., 2024). This distinction matters because a charge point can appear available in a system but still fail to deliver successful user-experience if authorization, payment, communication, connector availability, or power transfer fails. Uptime should therefore not be understood only as a technical status label, but also as part of the broader charging experience.

Reliability also has an equity dimension. Evidence from the United States suggests that disadvantaged communities may face fewer public chargers and more reliability issues, which means unreliable infrastructure can create unequal barriers for users with fewer charging alternatives (Yu et al., 2025). For GreenFlux, the implication is that technical charger signals should eventually be translated into service-impact indicators. CPOs need to understand not only which charger generated an error, but which issues are most likely to affect users, repeat across assets, or reduce charging success.

Technological

The technological pressure is twofold. Charging infrastructure is becoming more complex, while data-driven maintenance is becoming more feasible. Charging networks include different vendors, firmware versions, OCPP versions, charger types, connector configurations, and AC/DC systems. OCPP provides the communication layer between chargers and management systems, but it does not automatically provide complete diagnostic clarity. The Open Charge Alliance notes that OCPP can support downtime detection and, in some cases, downtime prevention, while emphasizing that OCPP itself is a message set between a charger and its management system (Open Charge Alliance [OCA], 2024).

This creates an important strategic distinction for GreenFlux. The value is not only in receiving OCPP messages, but in interpreting status notifications, error codes, vendor-specific messages, and connector-level event sequences into operationally meaningful signals. Newer protocol developments also point toward richer monitoring potential. OCPP 2.0.1 introduces a device model that can improve monitoring and remote-support capabilities, which the Open Charge Alliance links to reduced downtime and operating costs (OCA, 2024). For GreenFlux, this suggests a two-layer technology strategy: use

current OCPP 1.5/1.6 data more effectively now through issue families and health indicators, while preparing for richer OCPP 2.x device-level diagnostic later.

Data-driven fault detection also supports this direction, but it should not be overclaimed. Research on EV supply equipment monitoring suggests that anomalous charging-unit behavior can be detected from operational power profiles using autoencoder-based methods (Sakwa et al., 2024). This supports the technical feasibility of early fault detection, but operational predictive maintenance still depends on data quality, validated labels, feedback loops, and careful interpretation. For GreenFlux, the strongest pathway is therefore staged: signal interpretation first, diagnostic support second, charger-health indicators third, and predictive maintenance later.

Environmental

Reliable charging infrastructure supports the practical use of low-emission mobility. The environmental link is indirect but important. EV deployment contributes to transport electrification, but unreliable charging can weaken the practical attractiveness of EV adoption and reduce confidence in low-emission transport systems. As EV sales and public charging infrastructure expand, charging reliability becomes part of the operational foundation of the energy transition rather than a secondary technical detail (IEA, 2026).

For CPOs, this means maintaining charging infrastructure as dependable public infrastructure, not only as individual charging assets. For GreenFlux, the environmental implication should be framed carefully. Uptime features do not directly reduce emissions by themselves. Their contribution is enabling EV charging more usable, reliable, and scalable, which supports the broader transition toward electrified mobility. This framing connects the project's operational problem to the wider energy-transition agenda without overstating the environmental impact of software features.

Legal

The legal environment increasingly makes reliability measurable and reportable. In the United Kingdom, rapid public charge point networks must achieve 99% annual reliability measured as an average across each CPO's rapid network over the calendar year (Department for Transport, 2024). In the United States, the NEVI standard requires each charging port to have average annual uptime greater than 97% (Federal Highway Administration, 2023). These requirements do not directly define GreenFlux's full product strategy, they show that reliability is becoming a formal performance expectation in mature charging markets.

The legal and operational implication is that charger health and predictive maintenance features must be explainable, validated, and safe to use in decision-making. If alerts are unclear or unreliable, they can increase workload, create false urgency, or trigger wrong maintenance actions. This is especially relevant when predictive features are based on noisy logs, incomplete labels, or vendor-specific error

messages. For GreenFlux, the safer strategic position is to frame predictive maintenance as decision support rather than automated truth. Uptime reporting, issue prioritization, and diagnostic guidance should be traceable and interpretable before moving toward more advanced predictive models.

Based on above PESTEL analysis, the results highlight that the most relevant external pressures for this project are reliability accountability, CPO operational pressure, charger-data complexity, user-experience expectations, and the wider role of reliable charging in transport electrification. These pressures justify focusing the project on charger health and uptime intelligence rather than generic EV infrastructure deployment. For GreenFlux, the external environment creates a clear strategic direction by helping CPO customers turn fragmented charger data into structured signals for monitoring, triage, diagnosis, and future predictive-maintenance readiness.

Appendix 2. Stakeholder Landscape and Power-Influence Analysis for Charger Reliability

This appendix supports **Section 2.1** by mapping the stakeholders that influence charger reliability, uptime accountability, troubleshooting, and the adoption of maintenance-intelligence features. The analysis uses a power-interest logic inspired by stakeholder mapping approaches, where stakeholders are assessed based on their ability to influence an issue and their level of interest in the outcome (Mendelow, 1981). In this report, power/influence refers to a stakeholder's ability to shape decisions, data access, technical interpretation, maintenance action, regulatory expectations, or user perception. Interest refers to how strongly the stakeholder is affected by charger reliability and maintenance performance.

Charger reliability is treated as a multi-actor issue. Although CPOs are directly responsible for charger operations, reliability outcomes depend on software platforms, charger hardware, vendor diagnostics, field-service capacity, site conditions, grid or facility constraints, and user experience. EV charging literature shows that infrastructure development involves different stakeholder perspectives and priorities, including fast charging, smart charging, government intervention, interoperability, grid impact, and future energy integration (Wolbertus et al., 2020; van der Koogh et al., 2023; Karmaker et al., 2023). In the Dutch context, a dense and reliable public charging network is especially important because many households rely on public parking or parking garages (Netherlands Enterprise Agency [RVO], 2023).

For this report, the stakeholder landscape is grouped into four categories. First, **primary operational and customer-facing actors** are directly connected to charger uptime and troubleshooting outcomes. This includes CPO customers, GreenFlux support or customer success teams, and indirectly EV drivers.

Second, **internal capability actors** are GreenFlux teams that would translate maintenance-intelligence insights into platform capabilities. This includes the product/CPMS team and the data/machine learning team. Third, **Diagnostic and resolution actors** influence whether an issue can be technically interpreted, escalated, and resolved. This includes charger OEMs, vendor diagnostic systems, field-service partners, site hosts, and grid or facility operators. Fourth, **accountability and market pressure actors** shape external expectations around reliable charging services. This includes regulators, public authorities, electric mobility service providers (eMSPs), roaming platforms, and broader market actors. European policy increasingly emphasizes user-friendliness, interoperability, transparent information, and accessibility of public charging services (European Commission, n.d.).

Table 2.1 below highlights a high-level view of stakeholder influence and interest in charger reliability. It is used to distinguish which actors GreenFlux should manage

closely, keep satisfied, keep informed, or monitor when developing maintenance intelligence capabilities.

Table 2.1 High-level view of stakeholder influence and interest in charger reliability.

Quadrant	Stakeholders	Strategic reading for GreenFlux
Manager closely (high power/ high interest)	CPO customers; GreenFlux product team; GreenFlux support/customer success; GreenFlux data/machine learning team; key charger OEMs/vendors	These stakeholders are central to implementation and adoption. They influence whether maintenance intelligence becomes useful, trusted, and embedded in operational workflows.
Keep satisfied (High power/ medium interest)	Regulators/public authorities; large site hosts; grid or facility operators; major platform or roaming ecosystem actors	These stakeholders may not use the feature daily, but they shape reliability expectations, site constraints, data requirements, or market access conditions.
Keep informed (Medium power/ high interest)	EV drivers; field-service technicians; eMSPs; local site operators	These stakeholders experience or act on reliability problems, but often depend on CPOs, platforms, vendors, or field teams to change the system. Their feedback is important for usability and service improvement.
Monitor (Lower direct power/ lower direct operational interest)	Wider public; non-EV users; peripheral market observers	These stakeholders are not direct users of maintenance intelligence, but broader public perception may still affect the credibility of charging as part of the energy transition.

The table shows that CPOs, GreenFlux internal teams, and charger OEMs/vendors should be managed closely because they directly shape adoption, interpretation, and operational use.

Furthermore, **table 2.2** outlines the stakeholder register that translates the power-interest matrix into practical implications for GreenFlux. It focuses on each stakeholder's role in charger reliability, the main adoption or coordination risk, and the design implication for GreenFlux's maintenance-intelligence layer.

Table 2.2 Compact stakeholder register.

Stakeholder	Power-interest position	Role and main risk in charger reliability	GreenFlux design implication
CPO customers	High power/ high interest	CPOs are accountable for uptime, service	Provide prioritized review, clear monitoring outputs, action cues,

Stakeholder	Power-interest position	Role and main risk in charger reliability	GreenFlux design implication
		quality, operational cost, and customer experience. The main risk is that maintenance intelligence may not be adopted if it does not fit CPO workflows.	and uptime-related insights that reduce operational burden.
GreenFlux product team	High internal power/ high interest	The product team translates reliability insights into platform features. The main risk is that analytical findings remain one-off insights rather than scalable product capability.	Develop maintenance intelligence as a modular CPMS feature layer rather than a one-time analysis.
GreenFlux data/ machine learning team	High internal power/ high interest	The data/machine learning team develops feature tables, models, monitoring logic, and future predictive-maintenance capability. The main risk is weak labels and limited model validation.	Start with reusable feature tables, structured feedback labels, and explainable prioritization before more automated prediction.
GreenFlux support/customer success	High internal influence/ high interest	Support teams translate platform signals into customer-facing guidance. The main risk is alert fatigue or unclear ownership if outputs are not explainable.	Provide reason codes, confidence boundaries, escalation logic, and workflow-compatible review outputs.
Charger OEMs/vendors	High influence/ Medium-high interest	Vendors control hardware design, firmware behavior, diagnostic messages, and repair interpretation. The main risk is that vendor-specific signals remain ambiguous.	Build evidence trails and vendor-diagnostic mapping while avoiding unsupported root-cause claims.

Stakeholder	Power-interest position	Role and main risk in charger reliability	GreenFlux design implication
Field-service partners	Medium-high influence/ high interest	Field service actors execute inspections, repairs, replacements, and on-site interventions. The main risk is poorly prioritized field visits or missing feedback after intervention.	Include escalation context and structured outcome labels, such as resolved remotely, field intervention, site issue or vendor issue.
Site hosts/ property managers	Medium influence/ medium interest	Site hosts influence physical access, cabinet access, parking conditions, opening hours, and local operating constraints. The main risk is misattributing site/access issue to charger hardware.	Separate site/access context from charger technical signals where possible.
Grid or facility operators	Medium influence/ medium interest	Grid or facility actors influence electrical connection quality, local constraints, voltage behavior, and power availability. The main risk is misclassifying power-quality or site-electrical problems as charger failure.	Treat voltage, grounding, and power-quality signals with site context and avoid premature hardware-failure interpretation.
EV drivers/ end users	Indirect influence/ high interest	EV drivers experience unavailable, unreliable, or difficult-to-use charging services. The main risk is declining trust in public charging and reduced confidence in EV adoption.	Use end-user charging experience as an indirect performance lens, even when the direct customer is the CPO.
Regulators/ public authorities	High influence/ medium-high interest	Regulators define expectations for uptime, accessibility, interoperability, pricing transparency, and service quality. The main risk is	Support CPOs with better reliability monitoring, transparent uptime reporting, and service-quality evidence.

Stakeholder	Power-interest position	Role and main risk in charger reliability	GreenFlux design implication
		increasing accountability without sufficient operational intelligence.	
eMSPs/ roaming platforms	Medium influence/ Medium-high interest	eMSPs and roaming actors provide access, payment, and customer-facing service layers. The main risk is that poor status or availability data affects user experience across platforms.	Improve data quality and reliability information that can support downstream user-facing services.

The register confirms that charger reliability is not controlled by one actor. It depends on coordination between software, hardware, support, field service, site conditions, regulation, and user experience.

The stakeholder landscape shows that charger reliability is distributed across hardware, software, site conditions, grid or facility context, field service, regulation, and user experience. GreenFlux is not positioned to replace CPOs, charger OEMs, vendors, or field-service actors. Its stronger strategic role is to act as a software and data coordination layer that helps stakeholders interpret operational signals earlier and respond more consistently.

This has three implications. First, maintenance intelligence should not be designed as a black-box fault-diagnosis tool. CPOs, support teams, vendors, and field-service partners need explainable outputs that clarify the observed signals, asset level, operational context, and recommended review path. Second, predictive-maintenance readiness depends on feedback loops. Without structured review outcomes from support teams, CPOs, vendors, and field-service partners, GreenFlux will lack the labels needed to improve future models. Third, the solution should be stakeholder-aware, it should reduce operational burden for CPOs, provide usable context for support teams, create evidence trails for vendors and field service, and indirectly improve the reliability experienced by EV drivers.

Overall, the stakeholder analysis shows that charger reliability is distributed across CPOs, software platforms, charger vendors, field-service partners, site and grid-related actors, regulators, and EV users. For GreenFlux, charger health and uptime intelligence should therefore be understood as a coordination capability that connects operational data with clearer issue visibility, stakeholder context, and more consistent maintenance response.

Appendix 3. Internal Analysis: VRIO and 7S Assessment of GreenFlux's Maintenance-Intelligence Readiness

This appendix supports **Section 2.2** by assessing GreenFlux's internal resources and capabilities in relation to charger health and uptime intelligence. The VRIO framework evaluates whether resources and capabilities are valuable, rare, difficult to imitate (inimitability), and organizationally supported. It is grounded in the resource-based view of the firm, which links competitive advantage to firm-specific resources and capabilities (Barney, 1991).

In this report, VRIO is used qualitatively rather than as a numerical scoring exercise. This is more appropriate because the objective is to assess strategic relevance and implementation readiness, not to create an artificial ranking. **Table 3.1** highlights the results of VRIO analysis. The VRIO analysis is complemented by a compact 7S assessment. While VRIO helps identify whether GreenFlux has relevant internal capabilities, 7S examines whether the organizational conditions around strategy, structure, systems, skills, staff, style, and shared values are aligned with the development of maintenance intelligence. **Table 3.2** summarizes the overview of 7S analysis.

The 7S framework is useful because organizational capability depends not only on formal structure, but also on the alignment of hard and soft organizational elements (Waterman et al., 1980). The analysis draws on GreenFlux's public platform description, internal project context, stakeholder discussions, and the dataset structure used in this study.

Table 3.1 The results of VRIO analysis.

Capability cluster	Qualitative VRIO verdict	Relevance to the project
Platform and data layer	**Competitive parity/foundation**. Valuable and organized, but not rare by itself. Mature CPMS providers also offer platform monitoring, APIs, reporting, and operational controls.	Gives GreenFlux the base to connect charger data with CPO workflows. It becomes strategically stronger only when combined with diagnostic interpretation and workflow integration.
Operational charger data access	**Valuable but not sufficient**. Useful, but only partly rare, because other OCPP-based CPMS providers can access similar data types.	Enables analysis of charger states, connector events, session outcomes, repeated signals, and error patterns. The advantage depends on interpretation, not access alone.

Capability cluster	Qualitative VRIO verdict	Relevance to the project
Interoperability and ecosystem reach	**Adjacent advantage**. Valuable and more durable at scale because roaming, RaaS, standardization experience, and partner relationships accumulate over time.	Support cross-vendor and cross-network reliability logic but does not directly create predictive-maintenance capability by itself.
Algorithmic and smart charging capability	**Adjacent capability**. Shows expertise with algorithm-supported operation, but maintenance prediction requires different evidence.	Useful for analytics maturity, but predictive maintenance needs fault labels, repair outcomes, diagnostic validation, and feedback loops.
Diagnostic and maintenance tooling	**Core opportunity/temporary advantage**. Directly relevant, but monitoring, diagnostics, self-healing, and AI-supported maintenance are increasingly claimed by CPMS competitors	GreenFlux can strengthen differentiation through vendor-aware diagnostics, issue-family logic, connector-level interpretation, alert prioritization, and validated maintenance support.
Workflow and organizational integration	**Critical condition**. Potentially rare if executed well, the "organization" dimension is decisive.	Determines whether insights become action across product, support, customer success, and CPO workflows. Without this, analytics may remain technically interesting but operationally weak.

The VRIO assessment suggests that GreenFlux has a strong platform and ecosystem foundation, but not an automatic sustained advantage in predictive maintenance. Its CPMS platform, operational data access, interoperability activities, smart charging algorithms, and support-related services create a credible base for charger health and uptime intelligence. However, several of these resources are either common among mature CPMS providers or only indirectly related to maintenance. Raw charger data, basic monitoring, and broad ecosystem reach are valuable, but they do not by themselves create a defensible advantage.

A distinction should therefore be made between GreenFlux's broader company-level strengths and the specific capability assessed in the project. Roaming scale, smart charging, interoperability, RaaS operations, and ecosystem partnership strengthen GreenFlux's platform position, but they do not automatically solve the maintenance-intelligence problem. The relevant question for this report is narrower: whether these capabilities can be organized around the interpretation of charger health signals and embedded into CPO troubleshooting workflows.

The most project-relevant opportunity lies in diagnostic interpretation and workflow integration. GreenFlux can create stronger differentiation if it develops an interpretation layer that turns fragmented OCPP logs, vendor-specific error messages, station and connector-level events, and charging session outcomes into operationally meaningful signals. This would move the capability beyond basic monitoring toward issue prioritization, diagnostic guidance, and validated maintenance support. In VRIO terms, this is where GreenFlux can potentially move from competitive parity or temporary advantage toward a more durable advantage.

Several constraints shape this opportunity. Charger data is heterogeneous across vendors, firmware versions, protocol versions, charger types, and connector levels. Not all failures are visible in platform logs, especially when the root cause sits in hardware, vehicle behavior, grid conditions, user actions, or field-maintenance processes. Predictive models also required validated outcomes such as confirmed incidents, repair records, support-ticket linkage, or operational feedback loops. Without these, predictive alerts risk becoming noisy or difficult to trust.

From a VRIO perspective, this points to a staged product direction. GreenFlux should not treat predictive maintenance as an immediate black-box capability. The stronger pathway is to first improve signal interpretation, issue-family classification, connector-level diagnostic, and alert prioritization. These capabilities can then support charger-health indicators and, over time, predictive-maintenance readiness once validation and feedback loops become stronger.

The VRIO analysis shows that GreenFlux has relevant platforms, data, interoperability, and operational capabilities, but these capabilities do not automatically become a scalable maintenance-intelligence feature. The 7S assessment therefore complements VRIO by examining whether the organizational conditions around strategy, structure, systems, skills, staff, style, and shared values are aligned with this direction.

Table 3.2 Overview of 7S analysis.

7S element	GreenFlux relevance	Readiness implication	Risk if not addressed
Strategy	Reliability pressure creates a strategic fit for moving from basic monitoring toward charger health and uptime intelligence.	Maintenance intelligence can strengthen GreenFlux's CPMS value proposition if positioned as operational value for CPOs.	The initiative may remain a broad "predictive maintenance" claim without a clear strategic focus.
Structure	Ownership is distributed across product, data/machine learning, support,	Cross-functional ownership is needed to connect analytics, product design,	Insights may remain siloed and fail to become a scalable product capability.

7S element	GreenFlux relevance	Readiness implication	Risk if not addressed
	and customer-facing teams, while product teams would likely hold the development decision.	customer workflow, and feedback collection.	
Systems	GreenFlux has access to OCPP logs, CDR/session data, vendor information, and station/connector identifiers, but structured feedback labels are not yet available.	Reusable feature tables, review outcomes, feedback labels, and workflow integration are needed.	Predictive models may be difficult to validate or improve because operational outcomes are not systematically captured.
Skills	Maintenance intelligence requires a hybrid skill set: OCPP knowledge, charger diagnostics, data/machine learning, product thinking, and support-process understanding.	GreenFlux needs to combine technical analytics with operational troubleshooting experience.	A technically strong model may not translate into useful maintenance guidance.
Staff	Support and customer-facing teams are close to operational issues, while product and data teams can build the feature layer.	Roles should be clarified for review ownership, feedback capture, model monitoring, and customer roll-out.	No team may consistently own the feedback loop needed to improve the systems.
Style	A cautious, explainable, human-in-the-loop approach fits the current maturity of the data.	Automation should be introduced gradually, starting with interpretation, prioritization, and review support.	Over-automation may reduce trust if model outputs are treated as confirmed diagnosis.
Shared values	GreenFlux's role in scalable e-mobility aligns with reliability, trust, and operational continuity.	Maintenance intelligence should be framed as improving dependable charging services, not only as an AI feature.	The initiative may be perceived as a technical experiment rather than a customer-value capability.

Overall, the 7S assessment shows that GreenFlux has a strong strategic fit for maintenance intelligence, but implementation readiness depends on organizational alignment. The main internal challenge is not only whether GreenFlux has relevant data, but whether product, data, support, and customer-facing teams can create reusable systems, structured feedback loops, and workflow-compatible outputs. This supports the view that charger health and uptime intelligence should be developed as a staged internal capability rather than as an immediate fully automated predictive-maintenance solution.

Taken together, the VRIO and 7S assessments suggest that GreenFlux has a strong internal basis for charger health and uptime intelligence, but the defensible advantage is not yet automatic. VRIO shows that the opportunity lies in transforming fragmented platform data into operational judgement, while 7S shows that implementation depends on cross-functional ownership, structured feedback labels, reusable models/systems, and workflow-compatible outputs. The internal readiness challenge is therefore not only technical, but also organizational.

Appendix 4. Competitor and Best-Practice Analysis

This appendix supports **Section 2.3** by providing the evidence base for the competitor and best-practice analysis. The analysis separates two comparison levels. The first level compares general CPMS positioning on how providers compete in the platform functionality, commercial operations, interoperability, and ecosystem coverage. The second level compares maintenance-intelligence maturity: how providers frame diagnostics, self-healing, predictive maintenance, operational actionability, and evidence strength. This distinction is necessary because a provider can be strong in roaming, billing, payment, or white-label operations without necessarily being strong in predictive maintenance.

The benchmark is based on publicly available company websites, product pages, blogs, and official materials. These sources are useful for understanding market positioning, but they have limitations. Public claims do not always disclose the underlying analytical method, model performance, operational validation, or actual customer outcomes. Therefore, the benchmark should be read as a positioning and best-practice analysis, not as a verified technical audit of each provider's system.

4.1 Benchmark criteria

The benchmark uses two sets of criteria to avoid mixing general CPMS competitiveness with maintenance-intelligence integrity. **Table 4.1** summarizes the benchmark by comparing provider types, typical strengths, maintenance-intelligence relevance, and limitations.

CPMS positioning criteria

Strategic positioning refers to the main market narrative used by the provider, such as enterprise CPMS, white-label platform, transaction platform, charging operations, or reliability-led platform. Commercial and ecosystem layer refers to billing, tariffs, payments, roaming, settlement, white-label ownership, customer support, and partner ecosystem. Operational control layer refers to monitoring, alerts, remote commands, firmware or configuration support, incident handling, and operational workflows. Hardware and protocol coverage refers to multi-OEM support, OCPP compatibility, AC/DC support, and charger-model integration discipline. Energy and grid capability refers to smart charging, load balancing, peak shaving, grid integration, and fleet or site energy optimization. Trust and extensibility refer to APIs, certifications, enterprise readiness, integrations, marketplace, procurement hygiene, and standardization credibility.

Maintenance intelligence criteria

Claimed maintenance maturity refers to whether the provider claims reactive monitoring, proactive alerts, self-healing, predictive maintenance, AI diagnostics, or automated operational. Likely analytical method refers to the visible or implied logic

behind the feature, such as rule-based alerts, event-sequence logic, anomaly detection, machine learning, AI agents, or automated recovery rules. Operational actionability assesses whether the feature only detects problems or also helps diagnose, route, resolve, automate, or document the issue. Operational granularity assesses whether the feature works at network, station, connector, component, charger-model, or event-sequence level. Promised business outcome refers to the value promised by the provider, such as higher uptime, lower maintenance cost, faster resolution, fewer truck rolls, or better charging success. Evidence strength refers to the credibility of the claim based on public evidence, such as product description, case studies, quantified results, technical explanation, or only broad marketing language.

4.2 Benchmark synthesis

Table 4.1 Benchmark synthesis of CPMS and maintenance-intelligence providers.

Provider type	Typical strength	Maintenance-intelligence relevance	Limitation
Broad CPMS providers	Monitoring, billing, payments, roaming, white-label, APIs, and CPO operations.	Can embed reliability features directly into daily CPO workflows.	Maintenance claims may remain broad or depend on basic alerting rather than deeper diagnostic.
Enterprise CPMS/network platforms	Scale, multi-OEM operations, enterprise workflows, load management, and partner integrations.	Strong base for uptime reporting, operational control, and network-level reliability management.	May focus more on platform scale and operational control than predictive diagnostics.
Specialist AI/uptime platforms	Predictive maintenance, AI diagnostics, availability analytics, and issue prioritization.	Signal that maintenance intelligence is becoming a competitive frontier.	May depend on access to CPMS data, OCPP streams, charger logs, and integration into CPO workflows.
Hardware/OEM diagnostic tools	Deeper charger-level and component-level knowledge	Useful for model-specific or component-specific diagnostics.	Less neutral across multi-vendor CPO networks and may not integrate easily with CPMS workflow.
Internal CPO analytics teams	Direct access to internal operational priorities, support tickets, and maintenance outcomes.	Can tailor reliability logic to the CPO's own fleet and processes.	Harder to scale across customers and may require capabilities not all CPOs have.

The benchmark shows that basic CPMS capabilities are widely present. Providers compete through different emphases: Road.io is used here as a benchmark for modular white-label and transaction-oriented charging operations; AMPECO, Driivz, Monta, Chargecloud, Last Mile Solutions, and ChargePoint represent broad CPMS and operational-control platforms; EVailable.ai, Lemonflow, and Ampcontrol represent more specialized or adjacent maintenance-intelligence positioning. This distinction matters because a provider's strength in billing, roaming, or white-label operations does not automatically imply strong predictive-maintenance capability.

Maintenance intelligence appears to be an emerging frontier. AMPECO publicly describes remote management and maintenance features that include real-time alerts, fault-cause visibility, remote resolution, self-healing algorithms, and auto-fault recovery (AMPECO, 2026). Driivz frames proactive maintenance as a transition from reactive repairs to predictive performance with a workflow that includes detection, assessment, routing, resolution, and feedback collection (Driivz, 2025). Monta positions alert as real-time notifications for charge point issues, automatically fixing what it can and alerting the right people when it cannot (Monta, 2025). Virta explicitly claims self-healing functionality that can automatically resolve up to 30% of charger issues (Virta, 2024). EVailable.ai positions itself as an OCPP-based analytics, predictive-maintenance, and automation layer, while Ampcontrol positions 24/7 monitoring and predictive maintenance as ways to increase uptime and reduce operating costs (Ampcontrol, 2024; EVailable.ai, n.d.).

The key interpretation is that the competitive narrative is moving faster than the public evidence. Some providers appear to rely mainly on rule-based detection and auto-recovery logic, while others use broader AI or machine-learning language. The operational difference matters. A rule that restarts a charger after a known error is useful, but it is not the same as predicting a future hardware failure. Similarly, a dashboard that flags offline chargers improves visibility, but it is not the same as connector-level diagnostic or component-level failure prediction. For GreenFlux, the useful lesson is to avoid broad AI positioning unless the features can be explained, validated, and embedded into the troubleshooting workflow.

4.3 Porter Five Forces interpretation

Porter's Five Forces was used to interpret the competitive structure around CPMS and maintenance-intelligence features. The framework is useful because competition in this area is not limited to direct CPMS rivals. It also includes buyers, suppliers, substitute solutions, and new entrants (Porter, 2008).

Rivalry among existing competitors – High

Rivalry is high because many CPMS providers compete with overlapping platform capabilities, such as charge point monitoring, billing, payments, roaming, APIs, white-label applications, energy management, and operational control. As OCPP-based CPMS

functionality becomes more standardized, the baseline expectation of what a modern CPMS should do increases. This pushes competition toward reliability tooling, automation, analytics, and operational support.

Buyer power – High

Buyer power is high because many target customers are sophisticated CPOs, utilities, fleets, enterprise site operators, and eMSPs. These buyers increasingly care about uptime, multi-OEM support, compliance readiness, roaming coverage, integration quality, and operational efficiency. They can compare vendors not only on software features, but also on support quality, reliability claims, platform extensibility, and total operational burden.

Threat of substitutes – Moderate to high

The substitute threat is not only "another CPMS." Specialist uptime platforms, AI diagnostics providers, OEM diagnostic tools, operations partners, and internal analytics teams can solve parts of the maintenance problem outside the CPMS layer. EVailable.ai is an example of a specialist platform that positions itself as an upgrade to existing platform rather than a replacement CPMS (EVailable.ai, n.d.). This matters because GreenFlux may compete not only against full CPMS providers, but also against add-on analytics and reliability layers.

Threat of new entrants – Moderate

Software entry barriers are lower than hardware barriers, so new entrants can enter through niche analytics, dashboards, AI diagnostics, or predictive-maintenance tools. However, scaling from a niche tool into enterprise-grade CPMS operations is harder. It requires integration, certifications, operational credibility, customer trust, support capability, and computability with heterogenous charger fleets. This creates room for specialist entrants but also gives incumbent such as GreenFlux an advantage if they can embed maintenance intelligence into the existing platform.

Supplier power – Moderate

Supplier power is moderate because CPMS providers depend on charger OEM behavior, protocol implementation quality, firmware stability, network communication, and charger model differences. Open standards reduce dependence on any single hardware vendor, but real-world heterogeneity remains significant. This means uptime intelligence must be robust to messy data, inconsistent vendor messages, uneven firmware quality, and different charger behaviors across AC and DC assets.

Based on the analysis above, the competitor and best-practice analysis suggests that CPMS differentiation is moving from basic platform connectivity toward operational reliability and maintenance intelligence. However, the benchmark also shows that public claims remain uneven. "AI," "self-healing," and "predictive maintenance" are increasingly

used in market positioning, but public evidence often does not fully disclose the analytical logic, failure labels, validation process, or operational accuracy.

To avoid treating these claims as equivalent, the benchmark was interpreted using an explainability and validation lens. This distinction is important because predictive-maintenance systems are more useful when their outputs can be understood, validated, and acted upon by human users (Cummins et al., 2024; Amaliah et al., 2025). Trustworthy AI guidance also emphasizes validity, reliability, transparency, explainability, and accountability as conditions for responsible AI use (NIST, 2023). **Table 4.2** shows analysis on interpreting AI and predictive maintenance claims.

Table 4.2 Interpreting AI and predictive maintenance claims.

Public claim	What it may indicate	Evidence needed for stronger interpretation	Risk if unclear
Monitoring/ alerting	The platform detects status changes, faults, or availability problems	Data source, alert logic, entity level, alert frequency, and false-positive handling	More alerts without better prioritization
Diagnostics	The platform provides interpretation of possible issue type or cause	Error-code mapping, vendor diagnostic support, session context, and validation against resolved cases	Symptoms may be presented as causes
Self-healing/ remote recovery	The platform can trigger automated or semi-automated recovery actions	Scope of actions, success rate, rollback logic, human oversight, and safety boundaries	Automation may hide unresolved root problems
Predictive maintenance	The platform estimates future issues or failure risk	Prediction horizon, model type, training labels, validation methods, performance metrics, and operational feedback loop	“predictive” becomes a marketing label rather than a tested capability
AI-based prioritization	The platform rank assets or issue for review	Explainable reason codes, feature basis, calibration, lift/precision, test or holdout performance, and user workflow fit	Black-box scores may not be trusted or used

For GreenFlux, this creates both pressure and opportunity. Basic monitoring is unlikely to remain a strong differentiator, but GreenFlux can still build a credible reliability layer through vendor-aware diagnostics, connector-level issue logic, alert prioritization, session-impact indicators, and workflow-fit troubleshooting guidance. This would position GreenFlux not only as a CPMS that shows charger status, but as a platform that helps CPOs understand and act on charger health.

Appendix 5. OCPP Signals and Troubleshooting Flows

This appendix supports **Section 2.4** by providing technical and operational evidence behind the current-state operational analysis. It summarizes the main OCPP and CPMS data elements relevant to charger monitoring and maintenance-intelligence development and presents interview-based troubleshooting flows from CPMS/platform, DC charger owner/operator, and AC charger owner/operator perspectives. The purpose is not to provide a full technical explanation of OCPP or a complete process audit, but to clarify which platform signals are useful, where their limitations lie, and how these limitations shape real troubleshooting decisions.

5.1 OCPP 1.6-oriented CPMS data elements used in this project

Table 5.1 summarizes the OCPP 1.6-oriented data elements used in this analysis, including their maintenance value and limitations. It should be read as a project-specific interpretation guide for the available GreenFlux platform signals, not as a full comparison across all OCPP versions.

Table 5.1 OCPP 1.6-oriented data elements used for maintenance-intelligence interpretation

Data element	What is shows	Maintenance value	Limitation
Status notification	Reports charger or connector state changes, such as available, unavailable, faulted, charging, or suspended.	Useful for detecting availability changes, fault transitions, recovery patterns, and repeated abnormal states.	Does not always explain the root cause of the state change.
Error Code	Provides a standardized high-level error category.	Useful for first-level classification of abnormal events	Often too broad for diagnosis; categories such as Other Error may hide multiple issue types.
Vendor Error Code	Provides vendor-specific error detail.	Can help identify issue families, charger-model behavior, or recurring vendor specific pattern.	Not standardized across vendors and may require interpretation or validation.
Info field	Provides additional textual or contextual information attached to a status or error message.	Useful for interpreting ambiguous errors and grouping recurring messages.	Can be inconsistent, missing, unstructured, or difficult to compare across vendors.

Data element	What is shows	Maintenance value	Limitation
Transaction/ session records	Shows charging start, stop, duration, energy delivered, and session outcomes.	Help links technical events to user-facing charging success or failure.	A failed or short session may be caused by charger, vehicle, user, authorization, payment, or communication issue.
Station-level logs	Shows operational signals reported at charge point or station level.	Useful for identifying full-asset availability, communication, cabinet-level, or network-level issue.	May hide connector specific problems, especially in multi-connector or DC fast-charging assets.
Connector-level logs	Shows operational signals at connector or outlet level.	Important for identifying user-session impact, failures, and DC sequence-related issues.	More granular and therefore more complex to aggregate and interpret.
Firmware, protocol, and vendor meta data	Shows charger vendor, firmware version, protocol version, and technical context.	Helps compare behavior across vendors, firmware families, and OCPP versions.	May be incomplete, inconsistent, or insufficient to explain the issue without event context.
Remote command history	Shows whether reset, unlock, configuration, or other remote actions were attempted.	Useful for distinguishing remotely recoverable issues from cases requiring escalation.	Does not always explain why the remote action worked or failed.
Support or ticket information	Shows reported problems, customer complaints, escalation, or resolution notes if available.	Crucial for validating root cause and learning from resolved incidents.	May not be systematically linked to OCPP events or transaction records.

The main implication is that OCPP-based data is valuable but incomplete. It provides visibility into charger behavior, but not always direct root-cause diagnosis. For maintenance intelligence, the data must be transformed into operational context. This requires separating station and connector levels, compressing repeated messages, interpreting vendor-specific fields, linking events to charging-session outcomes, and validating patterns against support or field-resolution evidence where available.

OCPP 1.6 and OCPP 2.0.1: relevance for maintenance intelligence

The project analysis is primarily based on OCPP 1.6-oriented operational data, including status notifications, standard error codes, vendor-specific error information, transaction/session records, and station or connector-level identifiers. This fits the available dataset, but it should be distinguished from the richer monitoring capabilities available in OCPP 2.0.1. **Table 5.2** clarifies the current analytical scope and the longer-term relevance of newer protocol features for maintenance intelligence.

Table 5.2 OCPP 1.6 analytical scope and OCPP 2.0.1 roadmap relevance for maintenance intelligence.

Aspect	OCPP 1.6 relevance	OCPP 2.0.1 relevance
Main uptime-monitoring signal	Monitoring relies mainly on *Status* and *Error Code* from *Status Notification* messages.	Monitoring can use *Notify Event* messages and *Connector Status* from *Status Notification*, enabling richer event-based monitoring.
Diagnostic structure	Diagnostic meaning often depends on high-level error codes, vendor-specific fields, and surrounding event sequence.	The device model links events and errors to components and variables, making diagnostics more structured at component or logical-controller level.
Monitoring granularity	Useful for availability and fault-state monitoring but often limited for component-level interpretation.	Customizable monitoring can notify the CSMS when a component variable exceeds a threshold, supporting earlier warning signals.
Project relevance	Directly relevant to the current analysis because the dataset is structured on OCPP 1.6 style status, error, vendor, session, station, and connector signals.	Relevant for future-readiness because richer device-level monitoring may support more precise diagnostics and predictive-maintenance features.
Strategic implication for GreenFlux	GreenFlux can create short-term value by structuring and interpreting existing OCPP 1.6 signals more effectively.	GreenFlux can prepare a longer-term roadmap for richer OCPP 2.x diagnostics, while still needing validation, feedback loops, and workflow integration.

The comparison shows that OCPP 2.0.1 may improve the technical basis for device-level monitoring, but it does not remove the need for operational interpretation. OCA's uptime-monitoring guidance distinguishes OCPP 1.6 monitoring through *Status Notification*, *Status,* and *Error Code* from OCPP 2.0.1 monitoring through *Notify Event* and *Connector Status*; it also highlights the use of monitoring for crucial charging-station components in OCPP 2.01 (Open Charge Alliance [OCA], 2024). OCA material on OCPP 2.0.1 also describes the device model as a major improvement because events can be

linked to specific components and variables can be monitored against threshold (OCA, 2024). Therefore, the practical roadmap for GreenFlux is two-layered: extract more operational value from current OCPP 1.6 logs now, while preparing for OCPP 2.x-based diagnostic as charger networks and protocol adoption mature.

In addition, the comparison should not be read as a recommendation to immediately migrate the analysis to OCPP 2.0.1, because the relevance is strategic. It shows how richer device-model and event-monitoring capabilities could strengthen future maintenance intelligence if supported by charger implementation, CPMS integration, and consistent data quality.

Link to methodology

Above OCPP and CPMS data reality directly shapes the methodology used later in the report. Because raw status notifications do not equal confirmed failures, the analysis separates station and connector-level logs. Because repeated notification can inflate event counts, the analysis considers repeated-event compression and episode logic. Because charger behavior differs by asset type, the analysis separates AC and DC chargers. Because standard error codes can be too broad, the analysis examines vendor-specific error messages and hardware-oriented issue families. Finally, because charging success is user-facing, the analysis links operational events with transaction/session outcomes where possible.

The raw log line should therefore not be treated as the final analytical unit. A more useful unit is the interpreted operational episode. Like the affected asset, the entity level, the repeated or isolated nature of the signal, the likely issue family, the related session outcomes, and the available evidence of resolution. This is why the later analysis treats unhealthy charging as a data-driven proxy rather than a confirmed root-cause label.

5.2 Troubleshooting workflows and operational pain points across three perspectives

The following section summarizes three interview-based troubleshooting workflows from CPMS/platform, DC charger owner/operator, and AC charger owner/operator perspectives. The purpose is not to document every operational step in full detail but to identify recurring decision points and pain points in current troubleshooting practice. These flows help explain why charger-health intelligence should not be limited to issue detection, it also needs to support interpretation, routing prioritization, and feedback into future maintenance logic.

CPMS/platform perspective

As shown in **figure 5.1**, the CPMS/platform perspective highlights that the first visible signal of a charging problem does not always come from the backend. In some cases, the issue starts from the user side: a driver encounters an error, reports the location, charge point ID, or timestamp, and the platform team (most of them are through United Call Center (UCC)) then checks the charge point status and diagnostic log. This

matters because a user-facing charging failure may not map neatly onto a charger-side error code. A failed authorization, random card/token use, or driver-side issue may appear as a charging problem from the user perspective but sit outside the charger-failure scope.

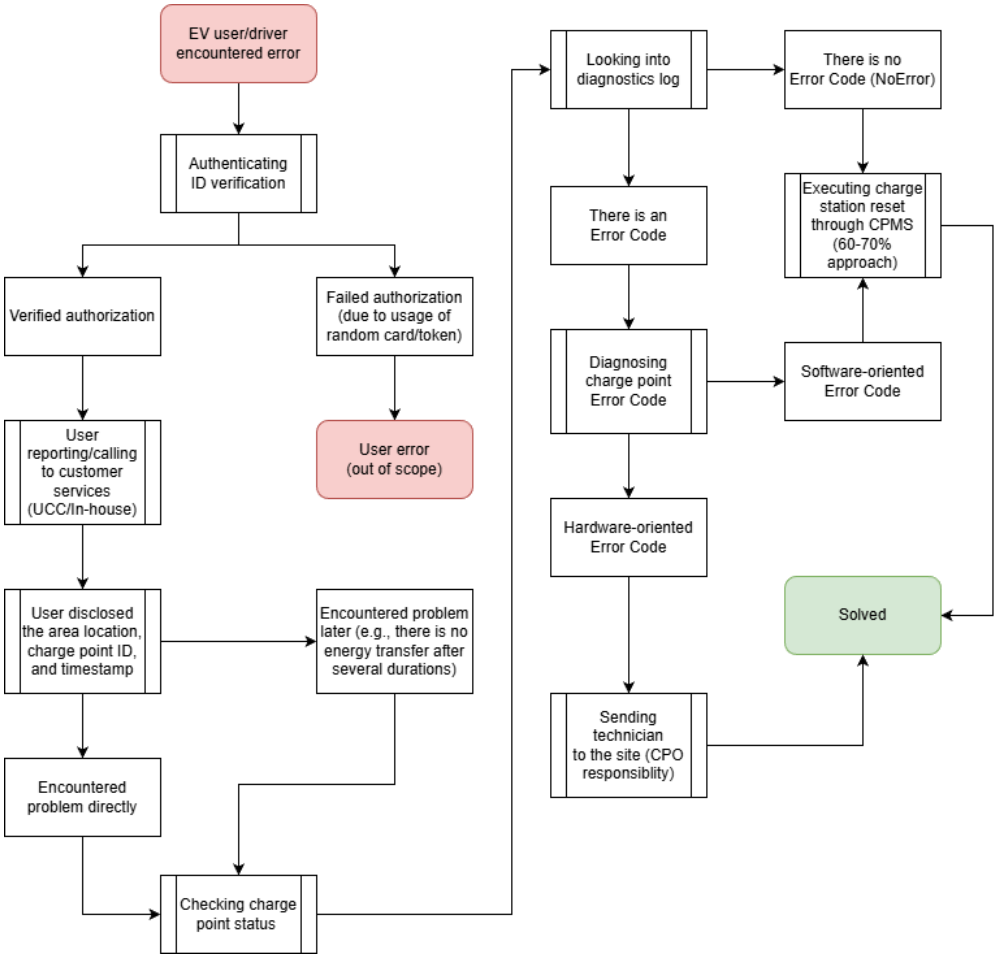

Figure 5.1 CPMS/platform troubleshooting flow based on interview mapping.

The pain point is that diagnostic logs do not always provide a clear interpretation. The flow shows that even when *No Error* is recorded, support may still attempt a reset or further check because the user experienced a problem. When an error code is present, the next step depends on whether it is software-oriented or hardware-oriented. Software-oriented issues may be handled remotely, while hardware-oriented issues often require CPO or technician intervention. This creates a boundary for GreenFlux: the platform can support detection and interpretation, but field resolution may sit outside the CPMS provider's direct control.

The implication for GreenFlux is that troubleshooting support should not treat error codes as the only source of truth. A more useful approach would connect diagnostic logs with session history, user reports, authorization context, and operational outcomes. This would help distinguish user-side or authorization issues from software-oriented and hardware-oriented charger problems, reducing unnecessary escalation and improving routing.

DC charger owner/operator perspective

The DC charger owner/operator perspective shows a more operationally urgent troubleshooting environment, as can be seen in **figure 5.2**. Because owner/operator treats DC charger as a cash cow for their EV charging business, so that it needs more attention. When an error occurs, the operator checks the EV portal and faulted charge station menu, then sorts the issue based on strategic location. This suggests that prioritization already exists, but it may be based mainly on location importance (e.g., busy highways) rather than a more systemic model of severity, recurrence, customer impact, revenue risk, or Service Level Agreement (SLA) relevance. The first response is often rebooting or restarting the charge station, which can solve a meaningful share of cases.

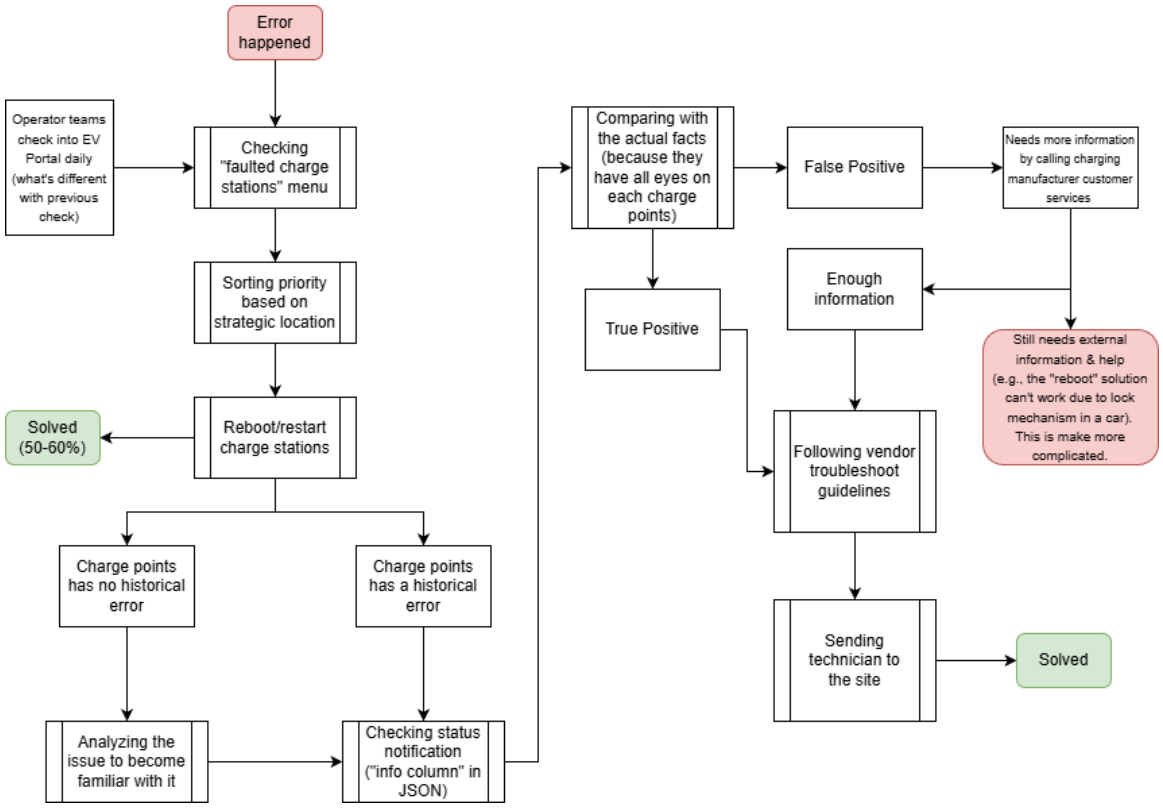

Figure 5.2 DC charger owner/operator troubleshooting flow based on interview mapping.

The main pain point appears after the basic recovery attempt fails. At that point, the operator needs to determine whether the issue is a false positive or true positive by checking historical errors, status notifications, and actual site facts. For instance, the operator received info in status notification JSON that stated, “cable theft,” but they are sure it has not happened, hence it is categorized as false positive. If there is enough information, the operator may follow vendor troubleshooting guidelines. If not, external information from the manufacturer or customer service may be needed. In unresolved cases, a technician may be sent to the site. This shows that DC troubleshooting often depends on vendor-specific guidance and real-world site validation, not only CPMS status information.

The implication for GreenFlux is that DC reliability support should focus on filtering and escalation logic. The platform can add value by helping operators decide whether a signal is likely to be real, recurring, and severe; whether reset or monitoring is sufficient; and when vendor or technician escalation is justified. Because DC chargers are more technically complex and often more business-critical, connector-level distinction, historical recurrence, and vendor-specific issue families are especially important.

AC charger owner/operator perspective

Figure 5.3 highlights the AC charger owner/operator perspective that the troubleshooting is often shaped by ambiguity rather than immediate technical severity. The process typically starts with an issue signal, customer complaint, failed session, or recurring abnormal behavior, followed by an initial support check and reset or restart. If the issue is resolved, the case may be closed or monitored for recurrence. If the issue persists or repeats, the operator checks which error appears and tries to interpret the likely issue family.

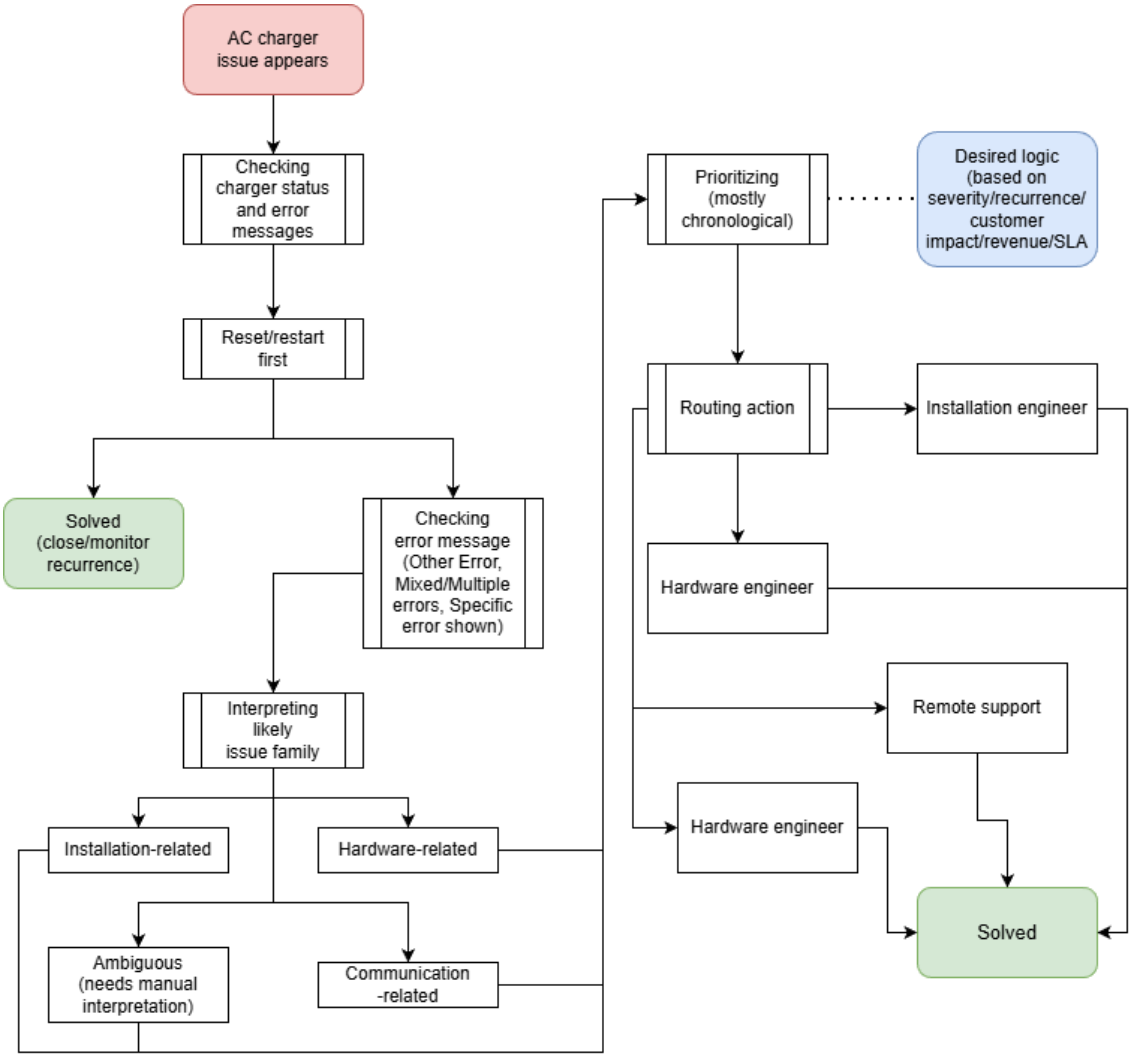

Figure 5.3 AC charger owner/operator troubleshooting flow based on interview mapping.

The main pain point is that AC error interpretation can be unclear. *Other Error* does not necessarily indicate charger hardware failure. In the interviewed case, it may refer to installation side-issues such as P1 meter, load balancing, or home electrical setup. Other cases may relate to RCD trip, ground fault, meter failure, communication issues, backend mismatch, or mixed error combinations. This makes routing difficult because the same high-level error category may require different actions: installation engineer, hardware engineer, remote support, or further investigation.

The AC flow also shows a prioritization gap. Cases may be handled chronologically rather than ranked by severity, recurrence, customer impact, revenue risk, or service level agreement (SLA) obligation. For GreenFlux, this suggests that AC maintenance intelligence should not start with a narrow hardware-failure prediction model. A more useful first step is issue-family classification, recurrence detection, and prioritization support. This would help CPOs distinguish installation-related, hardware-related, communication-related, and ambiguous cases more systematically.

Cross-perspective synthesis

Across the CPMS, DC, and AC perspectives, the recurring bottleneck is not the absence of data. It is the difficulty of turning available data into operational judgement. The CPMS perspective shows that user-facing issue may not map clearly to error codes. The DC perspective shows the need to separate true positive from false positive and to decide when reset, monitoring, vendor guidance, or technician dispatch is appropriate. The AC perspective shows how ambiguous and mixed errors complicate routing and prioritization.

These interview findings are consistent with broader reliability literature. Charging reliability should be understood as a process-chain issue, where failures may emerge from authorization, communication, payment, charger state, user interaction, power transfer, or maintenance response (Roeckle et al., 2022; CLEAResult, 2024). The ChargeX minimum required error-code work further supports the relevance of this finding. Inconsistent error reporting reduces the ability of operators to interpret, diagnose, and resolve charging problems consistently across vendors and systems (Savargaonkar et al., 2023).

These findings justify the methodological direction of the report. Because raw logs do not directly equal confirmed failures, the analysis separates station and connector-level signals, compresses repeated events, distinguishes AC and DC behavior, focuses on hardware-oriented issue families, links operational events with session outcomes, and treats unhealthy charging as a data-driven proxy than a confirmed root-cause diagnosis.

Appendix 6. Data Structure and Methodological Details

This appendix supports **Section 2.5** by providing additional detail on the data structure, method exploration, selected methodology, feature engineering, and baseline modelling setup. The purpose is not to document every code step, but to clarify how the analytical pipeline was constructed and why certain methods were selected or excluded.

Figure 6.1 provides an overview of the analytical workflow used in this study. The workflow starts from raw GreenFlux datasets, continues through data audit, entity-level table construction, error-code screening, issue-family mapping, feature engineering, train/test/keep-out splitting, baseline modeling, and recurrence benchmarking. The purpose of this figure is to show how raw operational data were transformed into evidence for issue-family prioritization and model-readiness assessment.

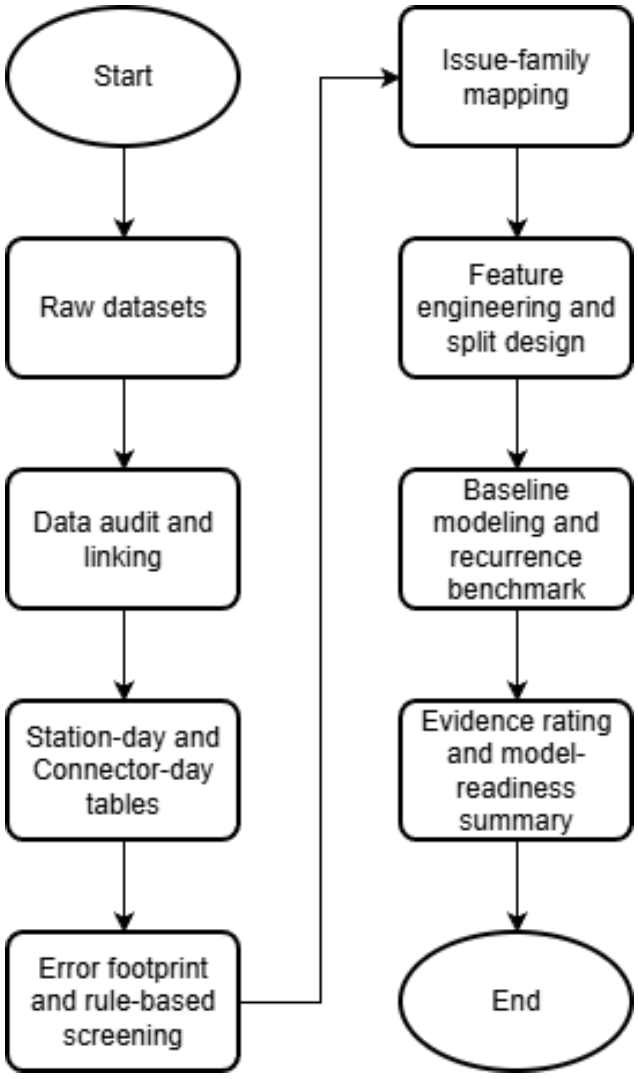

Figure 6.1 Overall analytical workflow from raw charger data to model-readiness assessment.

6.1 Dataset structure and analytical unit

The project data combines operational records with different temporal and entity structures. Status notifications are event-based records, CDRs are session-level

records, Meter Values are high-frequency telemetry records, and metadata provides asset-level context. To make these sources analytically comparable, the analysis aggregates relevant signals into daily analytical units at station and connector level. The observation period covers three months and is split into training (01/10/2025–30/11/2025), test (01/12/2025–15/12/2025), and keep-out (16/12/2025–31/12/2025) periods to separate model development from later validation. **Table 6.1** summarizes the main data sources, their analytical use, and their limitations.

Table 6.1 Data sources, analytical use, and limitations.

Data source	Raw structure	Analytical use	Main limitation
Station metadata	Station asset records	Vendor, protocol, firmware, age/context	Not causal by itself
Connector/ EVSE metadata	Static connector records	Connector-level entity mapping	Mapping can be complex
Location data	Site/location records	Spatial or operational context	Not always linked to cause
CDR/session records	Session-level records	Duration, energy, and session quality proxy	Failed sessions can have multiple causes
Station status notifications	Event-level logs	Station-level status and error behavior	May hide connector specific issue
Connector status notification	Event-level logs	Connector/session-facing issue behavior	More granular and noisy
Meter Values/ telemetry	High-frequency time-series records	Voltage/current context and validation	Uneven coverage and missingness

The station-day panel was used to evaluate station-level signals such as station-wide/network status changes, station-level error families, metadata, and daily session summaries. The connector-day panel was used to evaluate connector-specific patterns, especially for user/session-facing issues and DC charger behavior. This design allows the analysis to preserve the entity level of each signal instead of forcing all data into one generic charger-level unit.

6.2 Data cleaning and normalization rules

The cleaning process focused on making raw logs comparable while avoiding excessive assumptions. Status and error fields were normalized where needed, repeated notifications were compressed or interpreted carefully, and session records were filtered to reduce invalid or misleading observations. No Error, empty fields, repeated event bursts, very short sessions, zero-energy sessions, and missing telemetry were not treated as simple noise. Instead, they were reviewed according to their operational meaning.

The main cleaning logic followed four principles. First, invalid timestamps, impossible duration, and clearly inconsistent records should not drive issue-family conclusions. Second, repeated notifications should not be interpreted as independent failures without considering spam or persistence. Third, station-level and connector-level records should be kept separate because they represent different operational entities. Fourth, missing telemetry should not automatically remove an observation because missingness itself can reflect the limits of current observability.

6.3 Method exploration and selection logic

Predictive maintenance methods range from transparent rule-based monitoring to complex statistical, machine-learning, and deep-learning models. Reviews of predictive-maintenance methods show that there is no universally best approach. Method suitability depends on data structure, label quality, failure-event definition, sensor continuity, interpretability needs, and deployment constraints (Carvalho et al., 2019; Çinar et al., 2020). This project therefore treats method selection as a fit-for-purpose decision. Because confirmed repair labels and field outcomes are limited, the analysis prioritizes transparent screening, issue-family construction, and interpretable ranking over black-box prediction. **Table 6.2** summarizes the method families considered and the selection logic used in this study.

Table 6.2 Method families considered and selection rationale.

Method family	Why considered	Strength for this project	Limitation	Decision
Threshold/rule-based detection	Knowledge-based and rule-based monitoring remain useful when unhealthy conditions can be explicitly defined (Carvalho et al., 2019; Sharma et al., 2022)	Easy to explain; suitable for known operational conditions; useful for real-time monitoring	Thresholds can be arbitrary, charger-specific, and false-positive prone	Used for screening and diagnostic stress-testing, not final alerting
State-machine/ sequence rules	Event-based systems benefit from state/transition validation when operational states have clear meaning; this is suitable for OCPP-style status transitions	Captures abnormal transitions, persistence, and status-sequence logic	Requires strict definitions and charger/ Vendor-specific Interpretation	Used selectively faulted-link and sequence checks

Method family	Why considered	Strength for this project	Limitation	Decision
Weighted deterministic risk score	Risk scoring is operationally useful when calibrated probability is unavailable especially for triage and prioritization	Easy to communicate with operations; aligns with review queue and dispatch prioritization	Subjective weights: score is not real probability	Not selected as final analytical model; useful as possible product layer
Logistic regression/ Probabilistic classification	Classical machine learning is widely used in predictive maintenance; logistic regression provides an interpretable baseline for tabular binary classification (Carvalho et al., 2019)	Transparent, auditable, suitable for tabular features, and useful for baseline ranking	Needs calibration; assumes relatively simple relationships; sensitive to feature quality	Selected as baseline issue likelihood ranking model
Tree-based Machine Learning	Random forest, gradient boosting, and related methods are widely used in predictive-maintenance studies to capture non-linear patterns (Çinar et al., 2020)	Captures interactions and non-linear thresholds better than linear models	Less transparent; overfitting risk across heterogenous vendors and firmware versions	Future option after stronger validation and feedback labels
Anomaly detection	Useful in condition monitoring when labelled failures are scarce; anomaly detection can flag deviations from normal behavior (Sharma et al., 2022)	Can detect unknown or emerging abnormal behavior	Anomaly does not necessarily mean failure; hard to operationalize as maintenance action	Future exploration, not main method
Survival/RUL modelling	Suitable for time-to-failure and remaining-useful-life (RUL) problems when failure events and censoring are well defined (Jardine et al., 2006)	Strong for maintenance scheduling and time to-failure estimation	Requires clear failure-event definition, longer histories, and censoring structure	Not selected; labels proxy issue, not confirmed failures

Method family	Why considered	Strength for this project	Limitation	Decision
Deep learning/ Sequence learning	Deep learning and sequence models are used in mature PdM settings with larger sensor streams and labelled histories (Çinar et al., 2020)	Learns complex temporal and sensor patterns	Requires large labelled datasets, is harder to explain, and is risky without mature labels	Not suitable as main method now
AI copilot/ Diagnostic agent	Relevant as an interface layer for troubleshooting support, but reliable output depends on a validated diagnostic knowledge base and traceable evidence	Can help operators interpret logs and next actions	Risk of unsupported or overconfident advice if diagnostic logic is weak	Future layer after issue-family logic and validation mature

Condition-based and predictive-maintenance reviews note that ML-based maintenance systems benefit from richer monitoring data and interpretable outputs, but also face challenges in data quality, model interpretability, and deployment reliability (Sharma et al., 2022; Stiglic et al., 2020). For this project, the selected approach combines deterministic signal screening with probabilistic ranking. Rule-based screening clarifies raw signal behavior, issue-family mapping reduces vendor-specific error complexity, feature engineering converts raw events into interpretable indicators, and baseline logistic regression tests whether these indicators support issue-likelihood ranking on unseen periods. This staged design is more appropriate than direct black-box prediction because the current dataset lacks consistent field-confirmed failure labels.

The selection logic is important because the project objective is not to win a modeling benchmark, but to identify a feasible maintenance-intelligence pathway that GreenFlux can operationalize with its current data maturity. The available data are strong enough for issue-family interpretation, recurrence-based indicators, and explainable review prioritization, but not yet strong enough for autonomous predictive-maintenance claims.

6.4 Rule-based screening and faulted-link logic

Rule-based screening was used to test whether raw error signals were frequent, widespread, persistent, or linked to Faulted status. This step was necessary because raw frequency alone can be misleading. A high-frequency signal may represent a warning-like or contextual condition, while a lower-frequency signal may be more strongly associated with operational disruption. The screening logic therefore helps distinguish between broad signal footprint, recurring behavior, and closer fault-state linkage before moving to

issue-family interpretation and model-based ranking. **Table 6.3** summarizes the screening logic used in the analysis.

Table 6.3 Rule-based screening logic for raw error-signal interpretation.

Screening logic	Why it tests	Analytical use
Raw error frequency	How often error signal appears	Provides an initial signal footprint, but not severity.
Affected station/ connector share	How widespread the signal is across assets or connectors	Distinguishes network-wide patterns from localized issues.
Active-day persistence	Whether the signal repeats across days	Identifies recurring or unresolved issue behavior
Same-row Faulted overlap	Whether an error appears in the same log row as Status = Faulted	Provides a narrow proxy for direct fault-state linkage
Error-to-Faulted window	Whether an error is followed by Faulted within a selected time window	Captures warning-like signals that may precede faulted state
Station vs connector visibility	Whether the signal appears at station level, connector level, or both	Determines correct operational entity level

These rules were not treated as a final alert thresholds. They were used to classify signal behavior and identify which issue families deserved deeper feature testing. In particular, same-row Faulted overlaps and Error-to-Faulted windows helped distinguish warning-like or contextual signals from signals more closely linked to a faulted state. This matters because OCPP-style error reporting should be read as operational signaling, not as confirmed root-cause diagnosis.

6.5 Issue-family mapping

Issue-family mapping was used to reduce raw error complexity while preserving operational meaning. This step was necessary because OCPP standard error categories can be broad, vendor-specific fields are not standardized, and textual Info fields often contain diagnostic context that is lost when only raw Error Code is counted. ChargeX's minimum required error-code work emphasized that inconsistent error reporting is a barrier to diagnosis and reliability improvement, supporting the need for more structured interpretation of error signals (Savargaonkar et al., 2023). **Table 6.4** shows how raw signal sources were mapped into issue-family candidates.

The mapping was used to create issue-family flags rather than treating each raw error label as a separate modelling target. This is important because the same broad Error Code can have different operational meanings depending on vendor diagnostics, entity level, charger type, and surrounding event sequence. For example, a generic or broad error category may represent a warning-like signal in one context but a more fault-linked

signal in another. Issue-family flags therefore help preserve diagnostic meaning while making fragmented OCPP and vendor-specific signals more comparable for feature testing, review prioritization, and future feedback-loop validation.

Table 6.4 Issue-family mapping logic for raw OCPP and vendor signals.

Raw signal source	Mapping logic	Example issue family output
Error Code	Standard OCPP error category	Grounding, connector lock, power meter, voltage/grid
Vendor Error Code	Vendor-specific diagnostic detail	P1/Modbus, power-module, metering/pairing
Info field	Textual diagnostic context	Door opened/closed, No Free Power, sequence error
Status sequence	Operational state transition	Stop-loop, session-start, sequence/controller state
CDR/session outcome	User-facing session impact	Short session, failed/low-energy session
Telemetry context	Voltage/current condition where available	Voltage/grid validation

6.6 Feature engineering and feature-role classification

Feature engineering converted raw issue-family signals into candidate indicators for ranking. The goal was not only to maximize model performance, but to create features that can be explained for product, operations, and customer-support teams.

In this study, feature engineering was therefore treated as a core methodological step rather than a minor preprocessing task. Prior methodological work describes feature engineering as the process of extracting useful representations from raw data for machine-learning models (Zheng & Casari, 2018), while predictive-modeling guidance emphasizes that model quality depends not only on algorithm choice, but also on meaningful predictor representation and feature selection (Kuhn & Johnson, 2019). In this project, raw event logs, session records, telemetry, and metadata were transformed into operationally interpretable indicators such as recurrence, session quality, exposure, telemetry context, diagnostic context, and issue-family flags.

Spearman rank correlation was used as part of feature screening to identify monotonic associations between candidate features and issue-family indicators. This was appropriate because operational charger features are often non-normal, sparse, or skewed, especially recurrence counts, event counts, and session-quality rates. The correlation results were interpreted together with descriptive differences, standardized mean differences, and operational reasoning, rather than as standalone proof of causality. **Table 6.5** summarizes the main feature groups, example features, and their intended analytical roles in the modelling pipeline.

Table 6.5 Feature groups and intended analytical roles.

Feature group	Examples	Intended role
Recurrence	Same issue appeared yesterday, rolling 7-day issue count	Early warning and ranking signal
Session quality	Short-session rate, zero/low-energy sessions, valid-session share	User-impact context
Exposure/control	Total sessions, valid sessions, active-connector days	Avoid misleading raw counts
Telemetry/context	Voltage/current summaries or availability	Technical validation where coverage exists
Metadata	Vendor, firmware, protocol, station age	Context and segmentation
Diagnostic context	Faulted status, same-day diagnostic events, raw vendor messages	Dashboard/playbook explanation

A key distinction was made between early-warning candidates and same-day diagnostic/context features. Recurrence, prior-window issue counts, session-quality degradation, and exposure-adjusted features are more suitable for ranking because they can be constructed before the review target. Same-day Faulted status or same-day diagnostic counts may be useful for explanation, but they can create leakage if used as predictors without temporal separation.

6.7 Baseline modelling setup

Baseline logistic regression was used to test whether engineered features could support issue-likelihood ranking for selected issue families. In simplified form, the model estimates the log-odds and issue-family day as a linear combination of engineered features:

$$\log\left(\frac{p_i}{(1-p_i)}\right) = \beta_0 + \beta_1 x_{1i} + \beta_2 x_{2i} + \cdots + \beta_k x_{ki}$$

where p_i is the estimated issue-likelihood score for a station-day or connector day i, and $x_{1i}, \ldots, x_{ki}$ represent engineered features such as recurrence, session-quality indicators, exposure controls, telemetry context, metadata, or diagnostic signals. Positive coefficients indicate features associated with higher issue-family likelihood, while negative coefficients indicate features associated with lower likelihood. However, the coefficient should be interpreted as an issue-likelihood ranking score, not as a calibrated probability of hardware failure.

Logistic regression was selected because it is transparent, relatively simple to explain, and appropriate as a first benchmark for tabular operational data. Interpretable models are especially important when intended users need to understand why a case is prioritized and whether the ranking should be trusted (Rudin, 2019; Stiglic et al., 2020).

The modelling setup used train, test, and keep out periods to avoid judging performance only on the training data. Because issue-family days are relatively rare,

evaluation focused on PR-AUC, lift, precision/recall, and keep out stability rather than accuracy alone. Precision-recall evaluation is more informative than ROC interpretation in imbalanced binary classification because it focuses on performance for the positive class (Saito & Rehmsmeier, 2015).

The model score was interpreted as a ranking score, not a calibrated probability. This means that a high score indicates that a station-day or connector-day should be reviewed earlier than others for a given issue family. It does not mean that the score can yet be read as an exact probability of hardware failure. Calibrated predictive alerts would require stronger validation, confirmed repair labels, support-ticket linkage, and feedback from field or remote-resolution outcomes.

6.8 Methodological limitations

The main limitation is the absence of consistent verified root-cause and repair-outcome labels. As a result, the analysis relies on issue-family proxies and operational indicators. This is sufficient for prioritization and exploratory ranking, but not for causal diagnosis or fully automated predictive alerting.

A second limitation is that charger's behavior is vendor- and model-specific. Findings from Alfen AC and Alpitronic DC are valuable as representative cases, but they should not be generalized to all charger brands without additional testing. A third limitation is telemetry coverage. Voltage and current signals are valuable for validation, but uneven coverage means they cannot be used as the only basis for ranking. Finally, recurrence can be a strong feature but may also reflect persistent unresolved problems rather than true early warning. For this reason, recurrence-based ranking should be paired with feedback loops and operational validation.

The methodological implication is that the current project supports explainable issue-likelihood ranking, not full autonomous predictive maintenance. The next methodological step is to connect ranked issue-family outputs with ticket systems, remote command outcomes, vendor responses, and technician or installation-engineer findings. This would allow GreenFlux to move from proxy-based ranking toward calibrated prediction and validated maintenance playbooks.

Appendix 7. Detailed Analytical Evidence for Issue-Family Prioritization

This appendix provides detailed analytical evidence supporting the findings presented in **Section 2.6**. Its purpose is to show how raw charger logs, vendor specific diagnostic field, session records, telemetry context, and station/connector identifiers were transformed into issue-family evidence for maintenance-intelligence prioritization. The appendix therefore focuses on the analytical reasoning behind the selected AC and DC issue families, rather than repeating strategic discussion in the main text.

The results should be interpreted with clear boundaries. Error codes, vendor messages, recurrence patterns, and model scores are treated as operational signals, not as confirmed physical root-cause labels. The analytical objective is to assess whether selected signals can support explainable review prioritization, issue-family dashboarding, and future predictive-maintenance readiness. This means that the evidence is strongest when it shows stable ranking value, meaningful entity-level behavior, fault-linkage patterns, session-impact context, and operational interpretability.

The appendix is structured in five parts. **Section 7.1** explains the analytical setup, raw error-code behavior, and entity-level fault linkage. **Section 7.2** describes how raw OCPP and vendor signals were mapped into issue-family candidates while maintaining interpretation boundaries. **Section 7.3 and 7.4** present detailed analytical results for the representative AC and DC cases, Alfen and Alpitronic. **Section 7.5** summarizes the evidence strength and model-readiness level of each issue family, linking the analytical results to the recommended staged development of explainable uptime intelligence.

7.1 Analytical setup, raw error-code behavior, and entity-level fault linkage

The section explains the analytical setup used before issue-family mapping and model testing. The raw GreenFlux datasets contain different data structures: status notifications are event-based records, CDRs are session-level records, Meter Values are high-frequency telemetry records, and metadata provides asset-level context. To make these sources comparable, the analysis converted the data into daily analytical panels at station on one date, while a connector-day represents one connector entity on one date, defined by the combination of *ChargerStationId*, *EvseId*, and *ConnectorId*. This daily structure makes it possible to align status events, charging sessions, telemetry availability, and asset metadata before constructing issue indicators and model features.

The analysis used the training period to establish the raw signal baseline before applying filtered issue-family logic or model-based ranking. At station level, the training period covered 9,948 stations over 61 days, with 16,163 raw error events, 11,248 raw Faulted events, and 649,958 valid sessions. This equals approximately 2.49 raw error events and 1.73 raw Faulted events per 100 valid sessions. At connector level, the training period covered 8,801 stations, 30,316 connector entities, and 654,549 valid sessions. The connector-level lane reported a much larger raw burden, with 333,155 raw error

events and 158,823 raw Faulted events, equal to approximately 50.89 raw error events and 24.26 raw Faulted events per 100 valid sessions. Additionally, per 100 valid session normalization is sued to compare raw operational burden across entities with different charging activity levels, rather than treating high-usage stations or connectors as inherently more problematic. Valid session also defines as the charging session in which more than 5 minutes charging duration and more than 5 kWh energy uptake. **Figure 7.1** shows this difference in normalized raw signal burden.

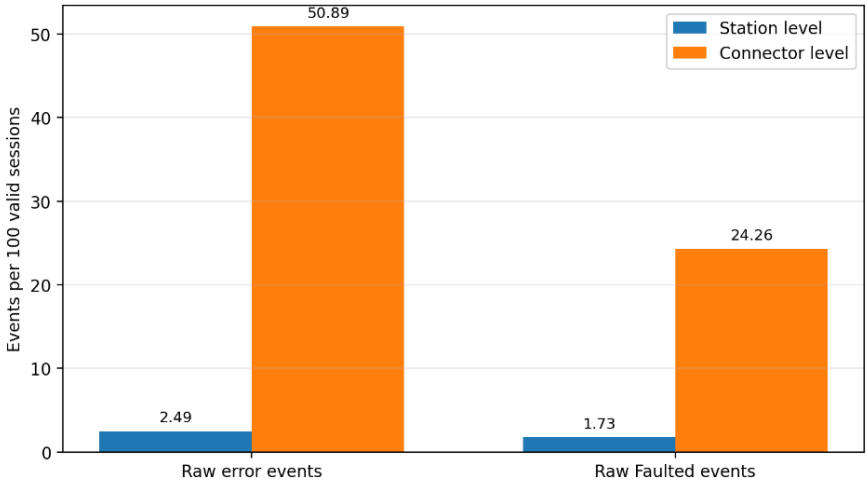

Figure 7.1 Raw error and Faulted event burden by entity level.

The difference in **Figure 7.1** should not be treated as evidence that connector-level operation is necessarily worse than station-level operation. Connector logs are more granular and more exposed to repeated notifications. The per-100-session metric is therefore used only as an exposure-normalized burden indicator, not as a measure of "errors inside sessions." This supports the need for filtering, entity-aware interpretation, and issue-family construction before treating raw logs as maintenance evidence.

The raw station-level error mix was dominated by Under Voltage, which contributed 56.64% of station-level raw error events, appeared across 1,278 stations, and showed a relatively high same-row Faulted overlap of 67.39%. At connector level, Undervoltage was also the dominant raw hardware-oriented signal, contributing 45.47% of connector-level raw error events, but its same-row Faulted overlap was only 4.91%. **Figure 7.2** compares this station-connector difference.

Figure 7.2 illustrates why raw frequency alone is insufficient for maintenance-intelligence interpretation. Under Voltage is frequent at both entity levels, but its same-row Faulted overlap differs sharply between station and connector logs. This suggests that connector-level Under Voltage may include many warning-like, repeated, or contextual notifications, while station-level Under Voltage appears more closely linked to a faulted state. Even so, both interpretations require additional context, such as voltage telemetry, recurrence, and session impact, before they can be treated as operationally meaningful issue candidates.

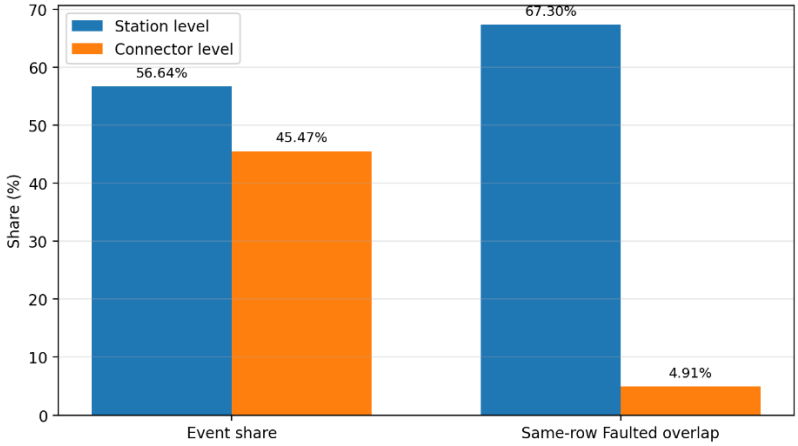

Figure 7.2 Under Voltage frequency and Faulted overlap at station and connector level.

Other hardware-oriented signals show the opposite pattern: lower event volume but stronger-fault linkage. At connector level, Ground Failure, Power Meter Failure, and Power Switch Failure had lower raw event volumes than Under Voltage, but very high same-row Faulted overlap, around 99.45%, 97.92%, and 9615%, respectively. Connector Lock Failure was also an important connector-specific candidate, with 27,110 raw events, 5,413 affected connector entities, and 69.53% same-row Faulted overlap. **Figure 7.3** summarizes the same-row Faulted overlap for selected connector-level signals.

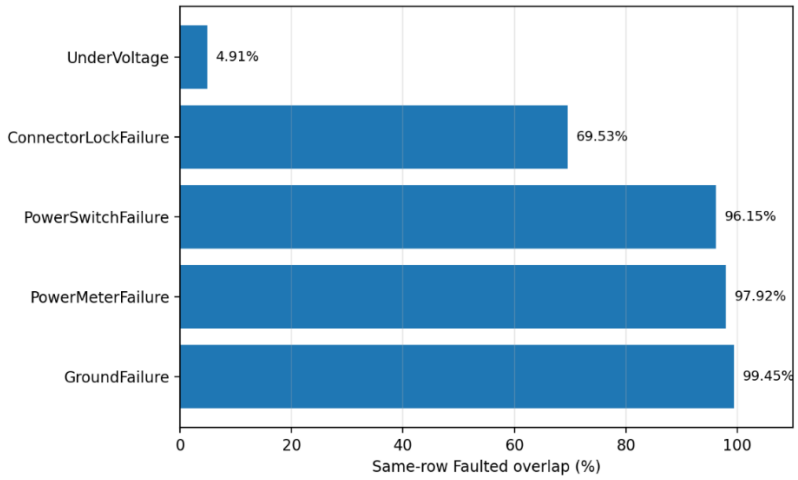

Figure 7.3 Connector-level same-row Faulted overlap for selected hardware-oriented signals.

Figure 7.3 reinforces the distinction between signal frequency and operational severity. A high-volume signal is not automatically the most maintenance-relevant, while a lower-volume signal may deserve attention if it is consistently linked to Faulted status. This is especially important for explainable prioritization because the objective is not to rank assets based only on how many raw events they generate, but to identify which signals are more likely to support meaningful review.

The station-connector comparison also shows that some hardware-oriented errors are much more visible at connector level than station level. This was especially clear for Power Meter Failure, where connector raw events were around 325 times higher

than station raw events, and for Over Current Failure, where connector raw events were around 455 times higher. Connector Lock Failure was also around 53 times more visible at connector level. **Figure 7.4** summarizes these visibility ratios.

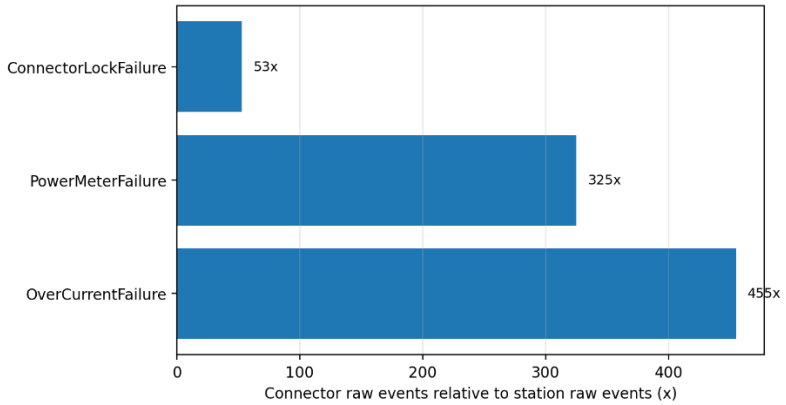

Figure 7.4 Station-connector visibility ratio for selected hardware-oriented error codes.

Figure 7.4 supports the methodological decision to keep station-level and connector-level analysis separate. Aggregating both levels too early into one flat charger-health score would risk hiding where an issue actually appears. For maintenance intelligence, this distinction matters because station-level signals may reflect asset, cabinet, site, or grid-related context, while connector-level signals may be closer to session-facing interface behavior, connector pat, and charging interaction.

Overall, the raw signal analysis shows why raw error-code counts are useful as a first footprint but insufficient as maintenance -intelligence evidence. A frequent signal may be warning-like or repeated, while a less frequent signal may be more closely linked to a faulted state. Raw error codes were therefore treated as operational signals requiring further interpretation through entity level, fault linkage, recurrence, vendor diagnostics, and session context. This motivates the next step: issue-family mapping, where raw OCPP and vendor-specific signals are grouped into operationally meaningful candidates for feature testing and review prioritization.

7.2 Issue-family construction and diagnostic interpretation boundaries

Issue-family mapping was used to reduce raw error complexity while preserving operational meaning. Raw Error Code labels were not treated as direct modelling targets because the same broad OCPP category can have different operational meanings depending on vendor diagnostics, entity level, charger type, recurrence, and surrounding event sequence. This is especially important for broad categories such as Other Error, which can appear uninformative when read alone but may contain useful diagnostic meaning when combined with Vendor Error Code, Info, and station or connector context.

The mapping therefore combined standard OCPP fields with vendor-specific and operational context. The main inputs were Error Code, Vendor Error Code, Info, Status, Timestamp, AC/DC charger type, vendor/model context, station-level or connector level entity, recurrence, fault-linkage behavior, session context, and telemetry availability.

Figure 7.5 summarizes the conceptual pipeline from raw charger signals to issue-family flags and analytical use cases.

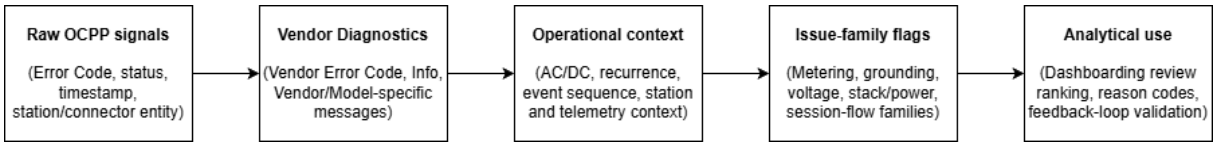

Figure 7.5 From raw charger signals to issue-family flags.

The issue-family flags were designed as operational review categories, not as confirmed physical root-cause labels. Their function is to make fragmented OCPP and vendor-specific signals more comparable, explainable, and suitable for future testing, review prioritization, reason-code explanation, and future feedback-loop validation. **Table 7.1** summarizes the signal dictionary used to construct the main issue-family candidates.

Table 7.1 Signal dictionary for issue-family construction.

Issue family	Standard OCPP anchor	Vendor/ Other Error support	Preferred analytical unit
Metering/P1/Modbus/ Measurement	Power Meter Failure	Other Error with metering, P1, metering-interface, or Modbus diagnostic meaning	Station-level and connector-level daily tables
Grounding/ electrical protection	Ground Failure	Other Error with grounding or protective-earth diagnostic meaning	Connector-level daily table first; station-level as context
Switching/component	Power Switch Failure	Other Error with switching or component diagnostic meaning	Connector-level daily table
Reader/authentication	Reader Failure	Other Error with reader or authentication diagnostic meaning	Station-level daily table
Load balancing/ smart charging control	Other Error	Other Error with load-balancing or smart-charging diagnostic meaning	Connector-level daily table; station telemetry as context
Voltage/grid quality	Under Voltage; Over Voltage	No main vendor-supporting field required	Station-level daily table first; connector level as context
Overcurrent/load protection	Over Current Failure	No main vendor-supporting field required	Connector-level daily table first; station telemetry as context

Issue family	Standard OCPP anchor	Vendor/ Other Error support	Preferred analytical unit
Connector lock/ Mechanical interface	Connector Lock Failure	No main vendor-supporting field required	Connector-level daily table
Firmware/ configuration	Other Error	Other Error with firmware or configuration diagnostic meaning	Station-level daily table
Communication/ backend	Other Error	Other Error with communication or backend diagnostic meaning	Station or connector level depending on diagnostic source
Thermal/temperature	High Temperature	No main vendor-supporting field required	Connector-level daily table

This mapping approach creates issue-family flags rather than treating each raw error label as a separate modeling target. The advantage is that it preserves diagnostic meaning while reducing vendor-specific fragmentation. For example, a broad Other Error signal can be grouped into metering, grounding, firmware, communication, or power-allocation families when the supporting vendor information and event context justify it. Conversely, when the supporting context is weak, the signal should remain dashboarding or exploratory rather than being used as a strong ranking target.

The interpretation boundary is important. Issue-family flags indicate that an asset or connector deserves review under a specific operational category: they do not prove that the underlying hardware component has failed. A metering/P1/Modbus flag does not automatically prove a meter defect, a grounding/electrical-protection flag does not automatically prove a broken power module. These interpretations require support-ticket linkage, vendor diagnosis, field-service feedback, or confirmed repair evidence.

The issue-family construction therefore supports the staged logic of the report. It enables dashboarding and explainable review prioritization today, while creating a more structured basis for future feedback labels and predictive-maintenance readiness. **Sections 7.3 and 7.4** apply this logic to the detailed Alfen AC and Alpitronic DC results.

7.3 Alfen AC detailed analytical results

This section applies the issue-family construction logic to the representative AC charger case, Alfen. The purpose is to identify which AC issue families are most suitable for explainable maintenance-intelligence prioritization. The analysis does not treat AC charger health as one generic score. Instead, it evaluates whether specific issue-families show enough raw traceability, fault-linkage behavior, feature association, validation performance, and operational interpretability to support dashboarding, review queues, or future predictive-maintenance development.

The Alfen AC results show that the strongest evidence is concentrated in a small number of issue families. Metering/P1/Modbus/measurement signals show the clearest readiness at station level and remain useful at connector level. Connector-level grounding/electrical protection also shows strong evidence because it combines raw diagnostic meaning, high Faulted linkage, and stable ranking value across validation periods. Other families, such as connector-level voltage/grid quality and switching/component, show moderate evidence. Connector lock and overcurrent/load protection are operationally relevant but should be treated more cautiously because their model evidence is weaker and more suitable for review support than automated alerting. **Table 7.2** summarizes AC issue-family model-readiness.

Table 7.2 AC issue-family model-readiness.

Issue family	Level	Test AP/PR-AUC	Keep-out AP/PR-AUC	Test lift	Keep-out lift	Interpretation
Metering/ P1/Modbus/ Measurement	Station	0.887	0.916	9.11	9.34	Best AC candidate for early issue-family dashboarding and review prioritization
Grounding/ electrical protection	Connector	0.674	0.668	8.60	8.47	Strong connector-level candidate with clear fault-linkage relevance
Metering/ P1/Modbus/ Measurement	Connector	0.533	0.611	8.23	8.19	Useful connector-level complement to station metering evidence
Voltage/grid quality	Connector	0.409	0.457	6.33	6.39	Useful for review prioritization, but requires voltage telemetry and site/grid context
Switching/ component	Connector	0.293	0.385	6.27	6.66	Operationally meaningful, but should remain human-reviewed
Voltage/grid quality	Station	0.124	0.047	5.24	5.00	High-volume signal but unstable ranking value at station level
Overcurrent/ load protection	Connector	0.073	0.060	6.50	6.19	Useful as a targeted review signal, not as an automated alert
Grounding/ Electrical protection	Station	0.062	0.123	6.67	9.06	Potentially relevant but too unstable for station-level use

Issue family	Level	Test AP/PR-AUC	Keep-out AP/PR-AUC	Test lift	Keep-out lift	Interpretation
Connector lock/ Mechanical interface	Connector	0.054	0.061	4.22	4.59	Operationally relevant connector signal, but limited model-readiness

Figure 7.6 visualizes the validation performance across AC issue families. Average Precision (AP) interpreted here as a precision-recall ranking metric, is used because issue-family days are highly imbalanced. In this setting, standard accuracy would be misleading because a model could appear accurate simply by predicting the majority non-issue class. AP is more useful because it evaluates whether issue-days are ranked above non-issue days.

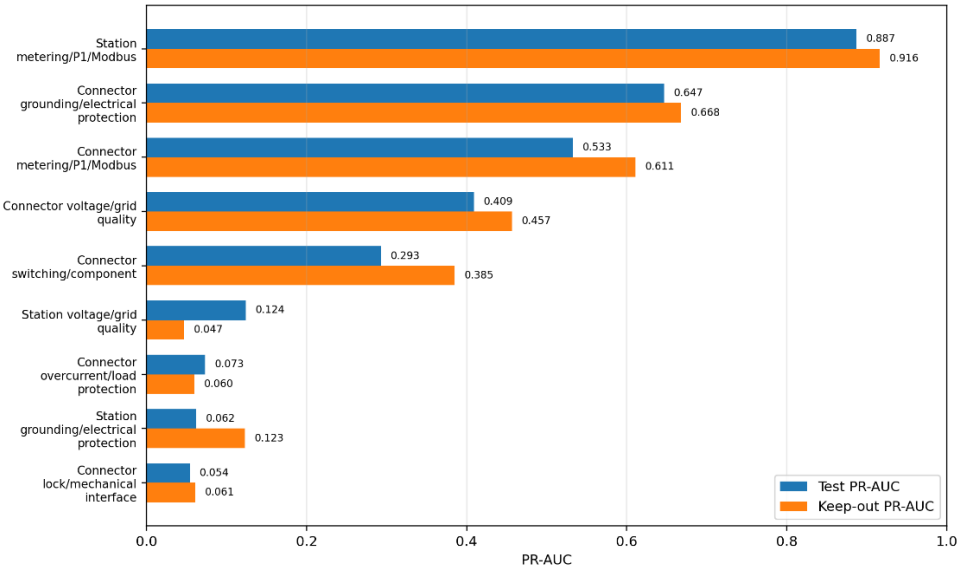

Figure 7.6 AC issue-family model-readiness by validation AP/PR-AUC.

Figure 7.6 shows a clear separation between strong, moderate, and weak candidates. Station-level metering/P1/Modbus/measurement is the strongest AC candidate, with high validation performance in both the test and keep-out periods. On the other side, Connector-level grounding/electrical protection and connector-level metering/P1/Modbus/measurement also shows strong evidence. This supports the view that AC maintenance intelligence should begin with specific issue-family ranking rather than a generic charger-health score.

To examine the two strongest AC candidates more deeply, **Figures 7.7** and **7.8** present precision-recall curves for station-level metering/P1/Modbus/measurement and connector-level grounding/electrical protection. These curves are important because they show how useful the ranking remains across different recall levels. In practical terms, they help to determine whether the model only identifies a few obvious cases or whether it can keep useful precision as more potential issue-days are reviewed.

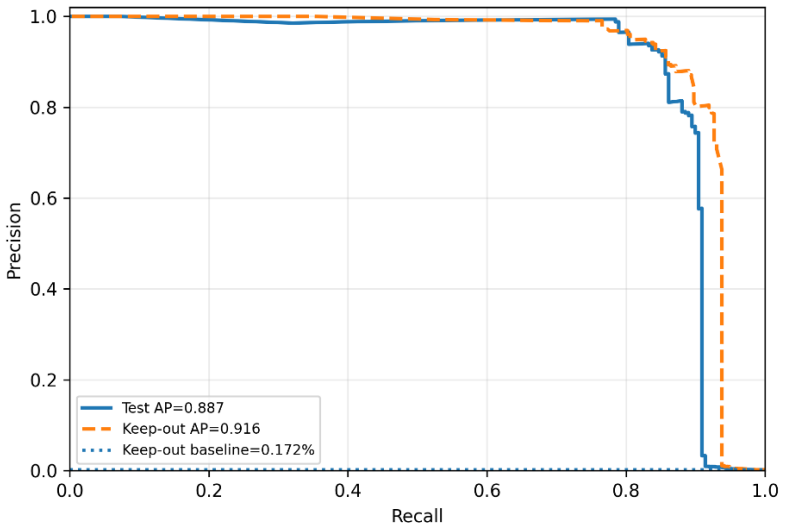

Figure 7.7 Precision-recall curve for station-level metering/P1/Modbus/measurement.

Figure 7.7 shows the strongest AC ranking result. The keep-out AP is 0.916, while the keep-out baseline is only 0.172%. The baseline represents the positive issue-family share in the keep-out period, meaning that a random review process would have a very low chance of finding true metering/P1/Modbus issue-days. Against this baseline, the model maintains very high precision across most of the recall range. This means that the top-ranked station-days are highly concentrated with likely metering/P1/Modbus issues.

The curve only drops sharply at very high recall, when the model is forced to include almost all possible issue-days and therefore also many lower-confidence cases. This is operationally important. It suggests that the model is most useful as a prioritized review queue: GreenFlux and CPO operational teams should focus first on the highest-ranked station-days, rather than trying to convert the full recall range into automatic alerts. For the first product release, this result strongly supports a station-level metering/P1/Modbus dashboard, reason-code explanation, and review-priority queue.

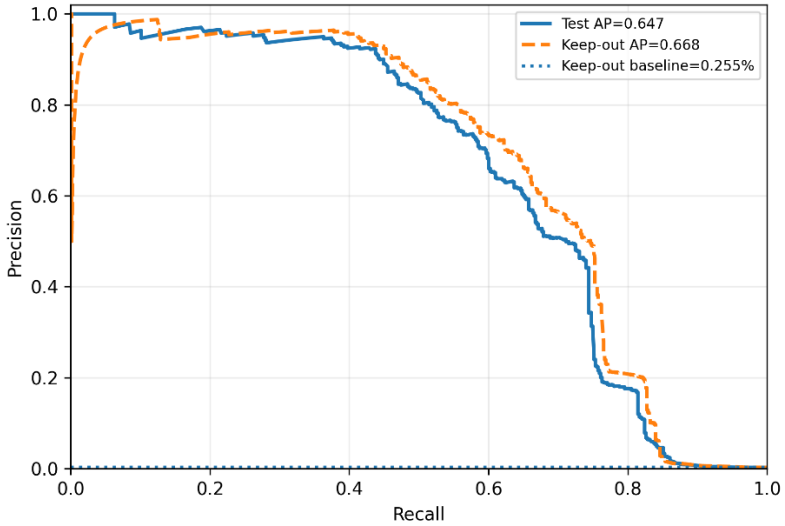

Figure 7.8 Precision-recall curve for connector-level grounding/electrical protection.

Figure 7.8 shows a different but still strong pattern. The connector-level grounding/electrical protection model has a keep-out AP 0f 0.668, compared with a keep-

out baseline of only 0.255%. This means that the model performs far above random ranking in a highly imbalanced setting. The curve remains substantially above the baseline across much of the recall range, although the decline in precision is more gradual than in the station-level metering case.

This pattern suggests that grounding/electrical-protection signals are operationally meaningful but less clearly separable than a station-level metering/P1/Modbus signals. The model can still rank likely issue-days well, especially in the higher-confidence part of the queue, but interpretation should remain human-in-the-loop. In operational terms, this issue family is suitable for connector-level review prioritization and troubleshooting support, not for autonomous diagnosis. A grounding/electrical-protection flag should prompt review of recurrence. Faulted linkage, connector-level status behavior, and field or installation context before escalation.

The feature-screening results also indicate that recurrence is an important part of the AC maintenance-intelligence logic. For metering/P1/Modbus/measurement, the strongest associated indicators include previous-day issue occurrence, seven-day recurrence, Faulted events, Available-to-Faulted transitions, and short-session behavior. For grounding/electrical protection, the strongest signals are also recurrence, Faulted linkage, and status-transition behavior. This suggests that repeated issue signals often carry more ranking value than isolated events.

However, recurrence should be interpreted carefully. It may indicate an unresolved operational issue, repeated warning-like notifications, or repeated exposure to the same site or connector condition. It should therefore support review prioritization rather than be treated as proof of an upcoming failure. Same-day Faulted events and status-transition features are useful for explanation and severity context, but they should not be overused as forward-looking predictors unless temporal separation is maintained.

The AC results therefore support a clear product implication for GreenFlux. The first AC maintenance-intelligence candidates should be station-level metering/P1/Modbus/measurement and connector-level grounding/electrical protection, supported by connector-level metering/P1/Modbus/measurement, voltage/grid quality, and switching/component signals as secondary review families. Connector lock and overcurrent/load protection should remain available as targeted operational review indicators. Overall, the Alfen AC evidence supports explainable issue family prioritization, not autonomous fault diagnosis.

7.4 Alpitronic DC detailed analytical results

This section applies the issue-family construction logic to the representative DC charger case, Alpitronic. The DC results should not be interpreted using the same logic as the AC results. For AC, the strongest candidates are mainly hardware-oriented and recurrence-driven issue families, such as metering/P1/Modbus and grounding/electrical protection. For DC Alpitronic, the evidence is more vendor-diagnostic-led and connector-

led. A large part of the diagnostic value is hidden inside Other Error, which means the standard OCPP error code alone is not sufficient for maintenance-intelligence interpretation.

The first important DC finding is that Other Error should not be treated as one generic or unknown error category. It behaves differently at station and connector level. At station level, Other Error is dominated by cabinet or enclosure-related signals, especially Door Opened and Door Closed. These signals are operationally useful because they indicate cabinet/site context, but they should not be directly interpreted as charger failure. At connector level, Other Error becomes more diagnostically diverse, covering session-flow, communication, metering-pairing, and stack/power-module availability signals. **Figure 7.9** shows this entity-level difference.

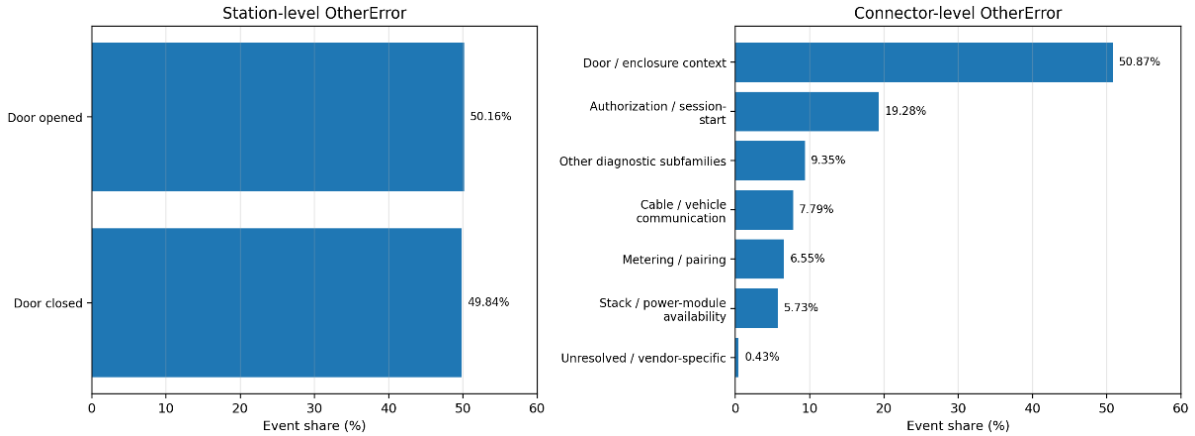

Figure 7.9 DC Other Error diagnostic-family composition by entity-level.

Figure 7.9 supports two interpretation points. First, station-level Other Error is mostly a cabinet-context container rather than a direct hardware-failure label. This is important because door opened/closed signals may reflect maintenance access, cabinet interaction, or site-level state changes. Second, connector-level Other Error contains several operationally meaningful subfamilies. Door/enclosure context is still the largest connector-level group, but authorization/session-start, cable/vehicle communication, metering/pairing, and stack/power-module availability are also visible. This confirms that Other Error should be decoded using entity level, vendor diagnostics, and operational context before it is used for prioritization. **Table 7.3** highlights diagnostic evidence and interpretation boundary for DC Alpitronic charger.

The second important DC finding is that model-readiness does not perfectly follow raw frequency or same-row Faulted linkage. **Figure 7.10** summarizes the validation performance across selected DC issue families using Average Precision (AP), interpreted as a precision-recall ranking metric. AP is used because issue-family days are imbalanced; the objective is to test whether likely issue-days are ranked above non-issue days, not whether the model can achieve high accuracy by predicting the majority class.

Table 7.3 DC diagnostic evidence and interpretation boundary.

Signal group	Key evidence	What it means	Interpretation boundary
Station-level door/ cabinet context	Station-level Other Error is dominated by Door Opened and Door Closed signals	Useful for cabinet/site context and avoiding false interpretation of door signals as charger failures	Should not be interpreted as charger failure without sequence, maintenance, or site context
Connector-level door/ enclosure context	Largest connector-level Other Error family by event share	Indicates that some connector-level Other Error volume may still be contextual	High volume does not automatically mean high severity
Authorization/ session-start	Large connector-level diagnostic family	Relevant for session-flow and charging-start experience	More relevant to user/session flow than direct hardware failure
Cable/vehicle communication	Visible connector-level diagnostic family	Indicates interaction between charger, cable, and EV communication	May involve charger-side, EV-side, or interaction context
Metering/pairing	Smaller event share but high same-row Faulted linkage in the diagnostic evidence	Operationally severe and important for troubleshooting	Strong fault linkage does not automatically prove a physical metering hardware defect
Stack/power-module availability	DC-specific connector-level diagnostic family	Indicates power allocation or power-module availability constraints	Weak same-row Faulted linkage does not mean low operational relevance
Unresolved/ vendor specific	Small but not-trivial residual diagnostic group	Captures remaining vendor-specific signals not cleanly mapped	Requires vendor or field-service interpretation before action

Figure 7.10 shows that the stack/power-module availability is the strongest DC ranking candidate, with a test AP of 0.623 and a keep-out AP of 0.589. Authorization/session-start and EV communication/session-context also shows moderate-to-strong ranking value, with keep-out AP values of 0.547 and 0.486. Grounding /electrical protection, safety/insulation/relay, emergency-stop/stop-loop, and metering/pairing show weaker but still operationally useful ranking value. This pattern supports a DC-specific interpretation: the most promising maintenance-

intelligence layer is not only a hardware-failure detector, but also a session-flow and power-allocation review layer. **Table 7.4** highlights the DC issue-family model-readiness summary.

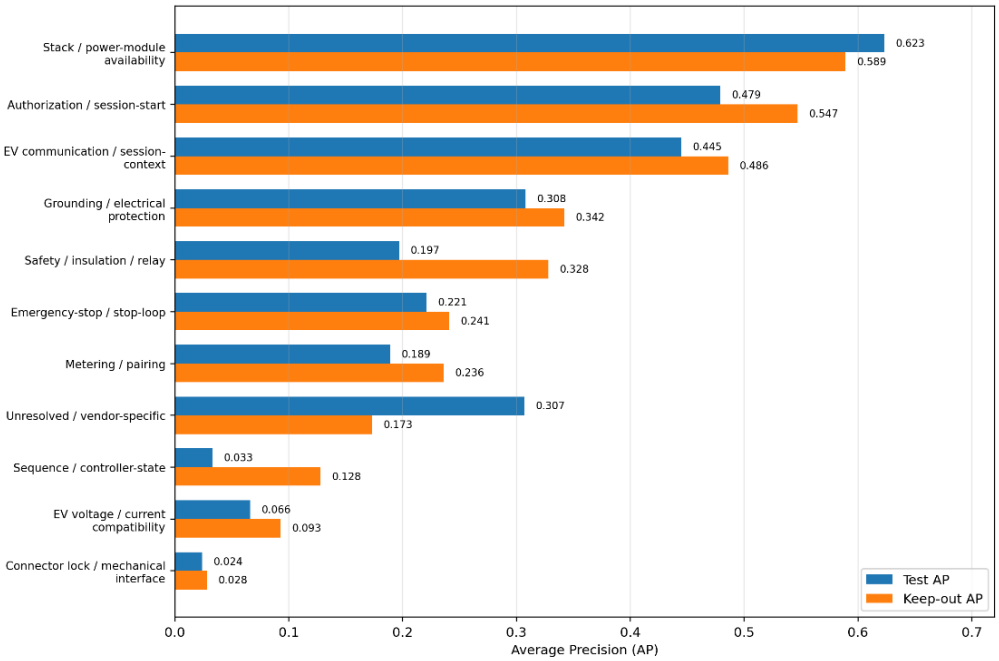

Figure 7.10 DC issue-family model-readiness by validation AP/PR-AUC.

Table 7.4 DC issue-family model-readiness summary.

Issue family	Test AP/PR-AUC	Keep-out AP/PR-AUC	Test lift	Keep-out lift	Interpretation
Stack/power module availability	0.623	0.589	7.37	6.80	Strongest DC candidate for connector-level review prioritization
Authorization/ session-start	0.479	0.547	4.60	3.91	Useful session-flow ranking candidate
EV communication/ session context	0.445	0.486	4.82	4.12	Useful for ranking connector-level communication/ session-context issues
Grounding/ Electrical protection	0.308	0.342	8.09	8.07	High lift, but likely sparse and party recurrence-driven
Safety/insulation/ relay	0.197	0.328	6.86	6.65	Operationally relevant but needs cautious validation
Emergency-stop/ Stop-loop	0.221	0.241	4.53	3.96	Useful as a review signal for stop-loop or interruption patterns

Issue family	Test AP/PR-AUC	Keep-out AP/PR-AUC	Test lift	Keep-out lift	Interpretation
Metering/pairing	0.189	0.236	7.68	8.04	Severe diagnostic family, but weaker ranking model than stack/power
Unresolved/ vendor specific	0.307	0.273	6.25	5.58	Needs vendor-specific interpretation before productization
EV voltage/current compatibility	0.066	0.093	4.32	3.87	Better treated as session-context review signal
Sequence/ controller-state	0.033	0.128	5.99	5.67	Potentially useful for diagnostic exploration
Connector lock/ mechanical interface	0.024	0.028	7.75	6.76	Operationally visible but limited model-readiness

The third DC finding is that the full model adds value beyond recurrence for several issue families, but not equally across all of them. **Figure 7.11** compares the keep-out AP and the full model with a recurrence-only benchmark. This is important because recurrence is often a strong operational signal in maintenance data. A model that only repeats recurrence logic may still be useful, but it should not be overclaimed as a more advanced predictive capability.

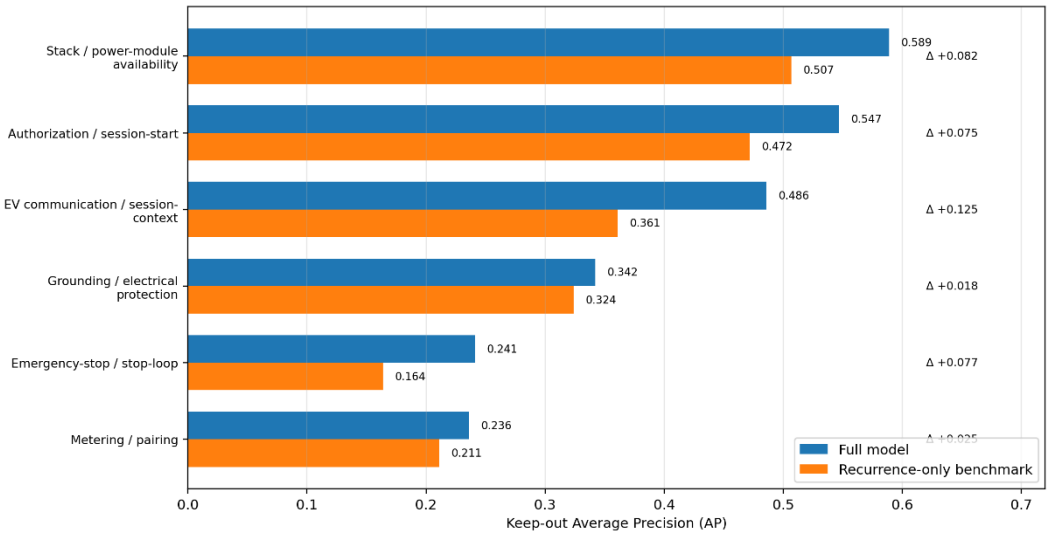

Figure 7.11 DC full model vs recurrence benchmark.

Figure 7.11 shows that the full model adds the largest value for EV communication/session-context, stack/power-module availability, authorization/session-start, and emergency-stop/stop-loop. The gain is smaller for

grounding/electrical protection and metering/pairing. This means that some DC families benefit from combining recurrence with additional session, status, and diagnostic features, while others remain closer to recurrence-driven ranking. For GreenFlux, this distinction matters because it helps decide which issue families are ready for model-assisted review queues and which should remain dashboard-only or expert-review categories.

Furthermore, the strongest individual DC deep-dive is stack/power-module availability. **Figure 7.12** shows the precision-recall curve for this issue family. The keep-out PR-AUC is 0.589, while the keep-out baseline is 6.5%. The baseline indicates the positive issue-family share in the keep-out period. A random review process would therefore be expected to identify true stack/power-module issue days at roughly the baseline rate. The model performs far above this baseline across most of the recall range.

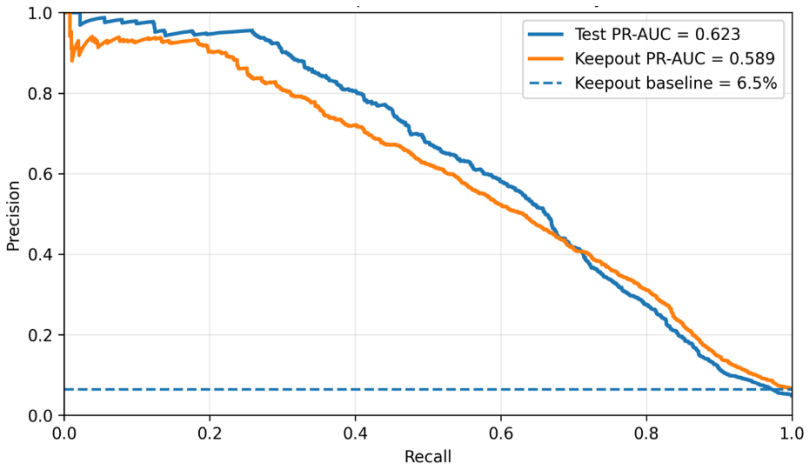

Figure 7.12 Precision-recall curve for DC stack/power-module availability.

Figure 7.12 shows that the model keeps high precision in the early and middle parts of the recall range, which is exactly the part most relevant for a review queue. This means that the highest-ranked connector-days are meaningfully enriched with likely stack/power-module availability issues. The curve declines as recall approaches 1.0, which is expected because the model must include lower-confidence cases to capture nearly all positives. Operationally, this supports stack/power-module availability as the first DC candidate for explainable connector-level prioritization. It should be framed as power-allocation or power-module availability review support, not as automatic confirmation that a power module has physically failed.

Figure 7.13 extends the deep-dive to two session-flow candidates: authorization/session-start and EV communication/session-context. These families are less hardware-specific than stack/power-module availability, but they are operationally important because they affect whether charging sessions start and proceed properly. The figure shows that both session-flow models rank substantially above their keep-out baselines. Authorization/session-start has a keep-out PR-AUC of 0.547 against a baseline of 15.5%, while EV communication/session-context has a keep-out PR-AUC of

0.487 against a baseline of 12.5%. The curves decline more gradually than the AC metering/P1/Modbus curve, which suggests that these DC issue families are broader and less sharply separable. However, their ranking value remains operationally meaningful, especially for prioritizing cases where session-start or EV communication issues may require review.

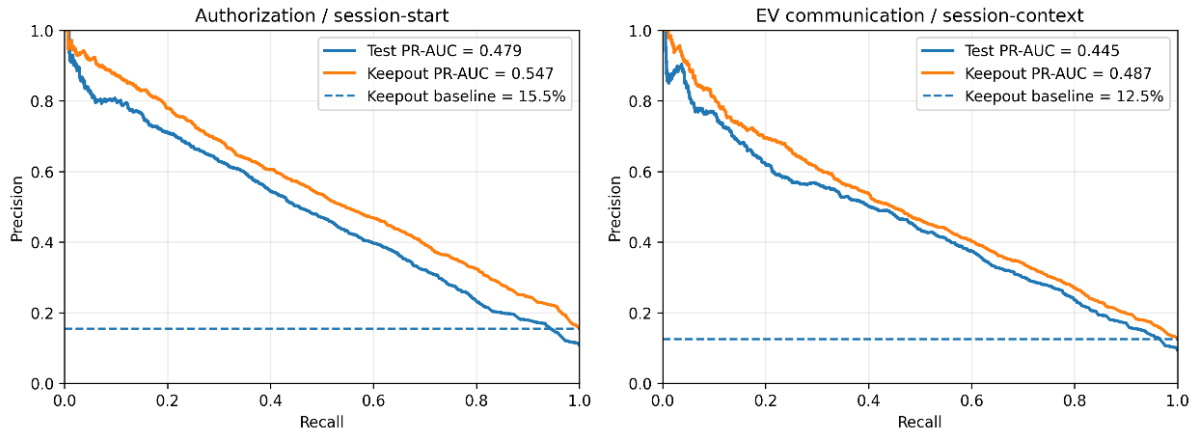

Figure 7.13 Precision-recall curves for DC session-flow issue ranking.

The practical implication is that DC maintenance intelligence should be designed around multiple diagnostic lanes. Station-level Other Error should mainly help prevent false interpretation of cabinet-door signals as charger failures. Connector-level Other Error should be decoded into issue families such as stack/power-module availability, authorization/session-start, EV communication/session-context, metering/pairing. And safety-related signals. Stack/power-module availability is the strongest DC ranking candidate, while authorization/session-start and EV communication/session-context are promising session-flow review candidates. Metering/pairing is operationally severe because of its fault-linkage evidence, but its current model-readiness is weaker and should therefore be treated as a dashboard and expert-review candidate before automated alerting.

Overall, the Alpitronic DC evidence supports a vendor-aware and connector-level maintenance-intelligence approach. The recommendation is not to create one generic DC health score, but to build issue-family-specific review queues that preserve Other Error diagnostic context, station/connector level, recurrence, session-flow behavior, and model-confidence boundaries. This is the most credible way to translate Alpitronic DC logs into explainable uptime intelligence without overclaiming automated fault diagnosis.

7.5 Evidence rating and model-readiness summary

This section consolidates the analytical evidence from **Section 7.1 – 7.4** into a model readiness summary. The objective is not to select one universal charger-health model, but to identify which issue families are currently suitable for dashboarding, explainable review prioritization, future machine-learning development, or only contextual monitoring. This distinction is important because the analysis shows that

issue families differ in signal clarity, validation stability, recurrence behavior, entity-level relevance, and operational interpretability.

The evidence rating used in this appendix combines four considerations: diagnostic interpretability, validation performance, operational relevance, and implementation readiness. A strong issue family does not mean that the model confirms a physical root cause. It means that the signals are sufficiently interpretable and stable to support an explainable maintenance-intelligence workflow. **Table 7.5** defines the evidence-rating logic used for the final synthesis.

Table 7.5 Evidence-rating criteria for issue-family model-readiness.

Evidence rating	Analytical meaning	Suitable use	Interpretation boundary
Strong	Clear issue-family meaning, stable validation performance, meaningful ranking value, and operationally interpretable signals	First pilot for dashboarding, review queues, and reason-code explanation	Support review prioritization, not automatic diagnosis
Moderate	Useful ranking or diagnostic value, but broader, less cleanly separable, or more dependent on contextual interpretation	Secondary review queue, human-in-the-loop troubleshooting support, and workflow validation	Needs additional validation before automated alerting
Weak but useful	Operationally meaningful signal, but weaker model-readiness, sparse positives, unstable validation, or limited separability	Dashboard visibility, targeted review, or diagnostic exploration	Should not be used as a standalone alerting target
Dashboard-only/ exploratory	Contextual, low-confidence, low-sample, or insufficiently mapped signal	Monitoring, descriptive analytics, and future evidence collection	Not suitable for ranking or prediction without stronger labels

Based on the logic, the strongest first candidates are AC station-level metering/P1/Modbus/measurement, AC connector-level grounding/electrical protection, AC connector-level metering/P1/Modbus/measurement, and DC connector-level stack/power-module availability. These issue families combine operational

meaning with stronger validation evidence and are therefore the most credible starting point for explainable uptime intelligence. Moderate candidates should be included only as secondary review families, while weak-but-useful candidates should remain visible for dashboarding and targeted investigations. **Table 7.6** summarizes the cross-case model-readiness assessment.

Table 7.6 Cross-case issue-family model-readiness summary.

Issue family	Entity level	Evidence rating	Recommended use	Next validation need
Metering/P1/ Modbus/ measurement	AC station-level	Strong	First pilot dashboard and station-level review queue	Link ranked cases to support tickets, metering diagnosis, and resolution outcomes
Grounding/ electrical protection	AC connector-level	Strong	First pilot connector-level review queue	Validate with field checks, installation context and confirmed electrical-protection findings
Metering/P1/ Modbus/ measurement	AC connector-level	Strong	Complementary connector-level review signal	Validate whether connector-level metering signals align with station-level metering context
Stack/power-module availability	DC connector-level	Strong	First DC connector-level review queue and power-allocation dashboard	Link with vendor diagnosis, power-module status, reset/reboot outcomes, and technician feedback
Authorization/ session-start	DC connector-level	Moderate	Secondary session-flow review queue	Validate whether ranked cases align with failed-start, user/session, or backend-authentication outcomes

Issue family	Entity level	Evidence rating	Recommended use	Next validation need
EV communication/ session context	DC connector-level	Moderate	Secondary session-flow review queue	Validate with session traces, EV-side context, cable interaction, and communication logs
Grounding/ Electrical protection	DC connector-level	Moderate	Human-reviewed safety/electrical-protection queue	Needs more positives, field validation, and vendor interpretation
Emergency-stop/stop-loop	DC connector-level	Moderate	Interpretation and stop-loop review signal	Validate with stop-event sequence, support tickets, and operational logic
Voltage/grid quality	AC connector-level	Moderate	Secondary review family with telemetry support	Requires stronger voltage telemetry and site/grid context
Switching/ component	AC connector-level	Moderate	Secondary connector-level review signal	Validates with component-level diagnosis and technician outcomes
Metering/pairing	DC connector-level	Weak but useful	Dashboard and expert-review candidate	Strong fault-linkage evidence should be connected to vendor diagnosis and confirmed outcomes
Safety/insulation/ relay	DC connector-level	Weak but useful	Dashboard and safety-oriented review candidate	Needs more labels, vendor confirmation and field-service feedback
Unresolved/ vendor-specific	DC connector-level	Weak but useful	Vendor-diagnostic exploration	Requires improved mapping of vendor-specific messages

Issue family	Entity level	Evidence rating	Recommended use	Next validation need
Connector lock/ mechanical interface	AC connector-level	Weak but useful	Targeted operational review indicator	Validate with connector-lock actions, unlock attempts, and user/session impact
Overcurrent/ load Protection	AC connector-level	Weak but useful	Dashboard and targeted review signal	Requires stronger electrical/load context and confirmed outcomes
Voltage/ grid quality	AC station-level	Weak	Contextual dashboard signal	Needs voltage telemetry, grid/site context, and separation from repeated warnings
Door opened/ door closed	DC station-level	Contextual	Cabinet/site context and false-failure prevention	Needs maintenance/site context; should not be used as fault diagnosis
Low-sample or unmapped vendor messages	AC/DC station or connector level	Exploratory	Monitoring and future mapping	Requires stronger diagnostic mapping and outcome labels

The summary indicates that GreenFlux should first operationalize issue-family ranking where three conditions overlap: the signal is interpretable, the model rank issue-days above baseline and the output can be translated into a clear review action. This is why AC metering/P1/Modbus, AC grounding/electrical protection, and DC stack/power-module availability are stronger first candidates than broader or more contextual families. At the same time, some lower-readiness families should not be discarded. They can still provide diagnostic value in dashboards, support investigations, and future feature development.

A second objective of this appendix is to provide a technical feature-engineering handover for GreenFlux's internal machine-learning development. The project created or identified several feature groups that are useful for issue-family ranking and future maintenance intelligence. These features should not all be treated equally. Some are already suitable for explainable review prioritization, while others are required before GreenFlux can move toward calibrated predictive alerts. **Table 7.7** summarizes the engineered feature groups and their recommended development priority.

Table 7.7 Engineered feature inventory for internal machine-learning development.

Feature group	Example engineered features	Why it matters	Recommended priority
Issue-family flags	Metering/P1/Modbus, grounding/electrical protection, stack/power-module, authorization/session-start, EV communication, connector lock	Converts fragmented OCPP and vendor diagnostics into interpretable ranking targets	Ready to use for explainable ranking
Recurrence features	Previous-day issue flag, seven-day issue count, rolling recurrence count, days since last issue	Captures repeated or unresolved issue behavior and is one of the strongest ranking signals	Read to use and refine
Event-intensity features	Total status events, error events, unique statutes, repeated events, state-change counts, state-change rate	Captures abnormal activity, unstable signaling, or repeated notification behavior	Read to use for ranking and reason codes
Status-transition features	Available-to-Faulted count, Error-to-Faulted window, Faulted overlap, Faulted duration on issue rows	Help distinguish warning-like signals from signals closer to operational disruption	Useful for explanation and severity context
Session-quality features	Short-session rate, failed-start/cancel pattern, valid-session share, session count, energy per session, duration per session	Links technical signals to charging outcomes and customer-facing impact	Ready to use, but should be validated with CPO workflow
Vendor diagnostic features	Vendor Error Code, Info keywords, Other Error subfamilies, stack/no-free-power-messages, P1/Modbus messages	Makes broad OCPP categories operationally meaningful and vendor-aware	High value, requires standardization
Entity-level features	Station-day, connector-day, station vs connector level, AC/DC, vendor/model, charger type, firmware	Prevents over-aggregation and supports charger-specific interpretation	Essential for all future models
Exposure-normalized features	Events per valid session, issue days per active connector, session-adjusted error rates, usage-adjusted burden	Reduces bias toward high-traffic assets and makes ranking fairer	Refine for production use

Feature group	Example engineered features	Why it matters	Recommended priority
Telemetry-context features	Voltage min/max/median, current, power, meter-value availability, telemetry gap rate	Needed for voltage/grid, metering, no-current-flow, and power-transfer interpretation	Improve coverage and data quality
Sequence features	Pre-issue state sequence, repeated status loops, transition patterns, ping-pong behavior, interruption/ anomaly sequence	Helps separate contextual noise from actionable issue progression	Worth developing for stronger early-warning logic
Remote-action features	Reset/reboot attempted, unlock connector, remote stop/start, action success/failure, time to recovery	Helps distinguish remotely recoverable issues from escalation cases	Required before operational alert automation
Support-ticket features	Ticket created, ticket category, issue owner, resolution text, time-to-resolution	Converts ranked signals into operational outcome labels	Required before calibrated predictive alerts
Field-service outcome features	Technician diagnosis, component replacement, installation issue, no fault found, repair confirmation	Needed to validate root cause and reduce false confidence	Required for diagnosis-oriented modeling
False-positive and review labels	Reviewed but not actionable, duplicate issue, automatically resolved, confirmed issue/no issue	Required to manage alert fatigue and calibrate thresholds	Required for alerting maturity

For near-term internal machine-learning development, GreenFlux should prioritize features that can support explainable ranking immediately: issue-family flags, recurrence, event intensity, status transitions, session-quality indicators, vendor diagnostic fields, and entity-level context. These features are already aligned with the current data structure and can support review queues with reason-code explanations. The next development layer should improve telemetry quality, sequence features, exposure normalization, and remote-action tracking. These features would make the ranking more operationally reliable and help distinguish persistent technical issues from contextual or high-usage noise.

However, the transition from ranking maintenance-intelligence to predictive alerts requires a different data layer. GreenFlux would need support-ticket linkage, remote-action outcomes, vendor diagnosis, technician findings, false-positive labels, and confirmed repair outcomes. Without these labels, a model can still rank likely issue-days, but it cannot reliably claim to predict physical failure, diagnose root cause, or

recommend field dispatch. **Table 7.8** summarizes the development of backlog needed to move from explainable ranking toward more mature predictive-maintenance capability.

Table 7.8 Future machine-learning development backlog.

Development need	Why it matters	Required data or process
Structured feedback loop	Converts review queue into learning data	Link ranked cases to support tickets, CPO feedback, remote actions, vendor diagnosis, and field-service outcomes
Ticket and resolution labels	Allows models to learn from operational outcomes	Ticket category, action taken, time-to-resolution, resolved/not resolved, repeated issue
Remote-action outcome tracking	Distinguishes remotely recoverable issues from escalation cases	Reset/reboot attempts, unlock actions, remote stop/start, success/failure, recovery time
Vendor-confirmed diagnostic labels	Improves vendor-aware issue-family mapping	OEM diagnostic categories, firmware fixes, hardware replacement, known error semantics
Field-service validation	Prevents overclaiming physical root cause	Technician findings, component replaced, installation issue, no fault found, repair confirmation
False-positive monitoring	Reduces alert fatigue and improves trust	Human review result, duplicate flag, non-actionable flag, automatically resolved cases
Calibration and threshold testing	Needed before probability-based alerts	Calibrated probabilities, precision targets, recall targets, top-K review performance, customer tolerance for false positives
Model monitoring and drift checks	Ensures performance remains reliable over time	Monthly/quarterly validation, vendor/model segmentation, charger firmware changes, data-quality checks
Product workflow validation	Ensures model output fits real operational work	Support-team review, CPO pilot feedback, escalation-path testing, reason-code wording validation

Taken together, the evidence-rating and feature-engineering synthesis supports the staged development path recommended in the main report. The current data is strong enough to support issue-family dashboarding and explainable review prioritization for selected AC and DC candidates. It is not yet strong enough for fully automated predictive maintenance or physical root-cause diagnosis. The most valuable next step for GreenFlux is therefore to operationalize a narrow set of high-evidence issue families, capture review and resolution outcomes, and use that feedback to improve model calibration, thresholds, and future alerting logic.

Appendix 8. Implementation Support for Explainable Uptime Intelligence

This appendix translates the analytical findings into implementation-support material for GreenFlux. While **Appendix 7** provides the detailed evidence behind issue-family prioritization, **Appendix 8** focuses on how these findings can be operationalized into an initial product and workflow concept. The aim is not to define a final software specification, but to outline what a first explainable uptime intelligence layer could look like, what feedback it would need, and what adoption risks should be controlled before GreenFlux moves toward more automated predictive-maintenance functions.

The proposed implementation logic follows the main recommendation of this report: GreenFlux should first develop human-in-the-loop decision support, not fully automated fault diagnosis. The strongest near-term opportunity is to help CPO customers interpret charger signals, prioritize which station or connector to review, understand why the case was ranked, and decide the next troubleshooting route. This approach keeps the output operationally useful while respecting the current evidence boundary: the available data can support issue-family ranking and review prioritization, but not yet confirmed physical root-cause diagnosis or autonomous field-dispatch decisions.

8.1 MVP product concept

The minimum viable product (MVP) should be positioned as an explainable uptime-intelligence layer on to of existing CPMS monitoring. Its purpose is to reduce the gap between raw charger logs and operational action. Instead of showing CPO users only status changes or error-code counts, the MPV should translate selected high-evidence signals into issue-family candidates, entity-level context, evidence explanation, and suggested review routes. The first version should focus on a narrow set of issue families with stronger evidence, such as AC metering/P1/Modbus/measurement, AC connector-level grounding/electrical protection, DC stack/power-module availability, and selected DC session-flow families.

The MVP should not replace existing station and connector views. It should act as a prioritization and explanation layer that helps users decide where to look first. This is important because charger data can be noisy, repeated, and vendor-specific. A generic charger-health score would risk hiding whether a signal appears at station-level, connector level, AC charger context, or DC charger context. The MVP should therefore preserve charger type, vendor, entity level, issue family, recurrence, session impact, and model-quality label. **Table 8.1** highlights operational questions and MVP outputs.

The first MVP should be deliberately narrow. It should start with issue-family dashboarding and review queue, then add reason-code explanation and feedback capture. This avoids two weak extremes: basic monitoring without interpretation, and premature AI automation without validated outcomes labels. The MVP should make the

system more useful to operators today while creating the feedback data needed for future machine-learning maturity. **Table 8.2** shows MVP capability stack and maturity status.

Table 8.1 Operational questions and MVP outputs.

Operational question	MVP output	Example implementation
What happened?	Issue-family classification	Metering/P1/Modbus, grounding/electrical protection, stack/power-module availability, authorization/session-start, EV communication/session context
Where did it happen?	Entity-level interpretation	Station-level queue, connector-level queue, or cross-level context
Does it matter?	Evidence and priority explanation	Recurrence, session impact, Faulted linkage, AP/model-readiness level, model quality label
Why is this case ranked?	Reason-code explanation	“Repeated in last 7 days,” “linked to short session,” “vendor diagnostic indicates stack/power availability,” “connector-level fault-link evidence”
What should be done next?	Suggested review route	Monitor, remote reset/reboot, check connector/session history, contact CPO support, vendor escalation, field inspection
What was the outcome?	Feedback-loop capture	Ticket result, remote-action outcome, vendor diagnosis, technician findings, false positive, confirmed repair

Table 8.2 MVP capability stack and maturity status.

MVP capability	Description	Current feasibility	Implementation note
Issue-family dashboard	Shows selected issue families by station/connector, vendor, charger type, recurrence, and session context	Feasible now	Start with high-evidence AC dan DC candidates; avoid showing all raw errors as equal
Station-level and connector-level review queue	Rank assets or connectors for human review based on issue-family evidence	Feasible as human-in-the-loop decision support	Keep station and connector queues separate to avoid over-aggregation
Reason-code explanation	Explains why a case is ranked or flagged	Feasible with careful wording	Use operational evidence such as recurrence, session impact, vendor diagnostics, and Faulted linkage

MVP capability	Description	Current feasibility	Implementation note
Model-quality labels	Distinguishes strong, moderate, weak-but-useful, and dashboard-only issue families	Feasible now	Prevent false confidence and makes evidence boundaries visible
Suggested troubleshooting route	Provides next-step guidance such as monitor, remote action, support review, vendor escalation or field inspection	Feasible as guided workflow, not automatic instruction	Should be validated with support teams and CPO users
Feedback capture	Links ranked cases to operational outcomes	Required next step	Needs ticket linkage, remote-action outcome, vendor diagnosis, technician feedback, and false-positive review
Calibrated predictive alerts	Converts issue-family ranking into probability-based or threshold-based alerts	Later-stage capability	Should only be considered after feedback labels, calibration, and false-positive tolerance are validated

In product terms, the MVP should be designed around "explainable prioritization." A CPO user should be able to open a ranked case and understand the entity, issue family, evidence, confidence boundary, and recommended review route. For example, an AC station-level metering case my be ranked because of repeated P1/Modbus-related diagnostics and short-session patterns, while a DC connector-level stack/power-module case may be ranked because of recurring vendor diagnostic related to power availability. In both cases, the output should guide review, not claim final diagnosis.

The MVP should also make uncertainty visible. Strong issue families can be used for the first review queues, moderate families can be included with cautious labels, and weak-but-useful families should remain dashboarded or manually reviewed. The design keeps the product credible and reduces the risk of alert fatigue. It also gives GreenFlux a practical route to learn from user behavior: which ranked cases were reviewed, which were ignored, which led to support action, and which were false positives.

A simple MVP workflow can therefore follow six steps: ingest charger signals, map them into issue families, rank station or connector cases, explain the reason for ranking, suggest a review route, and capture the outcome. This workflow preserves the analytical boundary of the project while turning the findings into a feasible product direction.

8.2 Feedback-loop and ownership design

The MVP described in **section 8.1** can only mature into predictive-maintenance capability if GreenFlux captures operational feedback from ranked cases. Without this feedback, the system can still support issue-family ranking and review prioritization, but it cannot reliably learn which signals are actionable, which are false positives, which can be resolved remotely, and which require vendor or field-service escalation. The feedback loop is therefore the missing layer between explainable prioritization and future calibrated predictive alerts.

Figure 8.1 illustrates the proposed learning loop. A ranked issue-family case is first reviewed by a human operator, support team, CPO operational team, or customer-facing team. The reviewer then decides the next action, such as monitoring, remote reset/reboot, vendor escalation, or field inspection. The outcome of that action should be captured and returned to the data pipeline. Over time, this creates better labels, improves reason-code quality, supports threshold tuning, and helps GreenFlux decide which issue families are ready for more automated alerting.

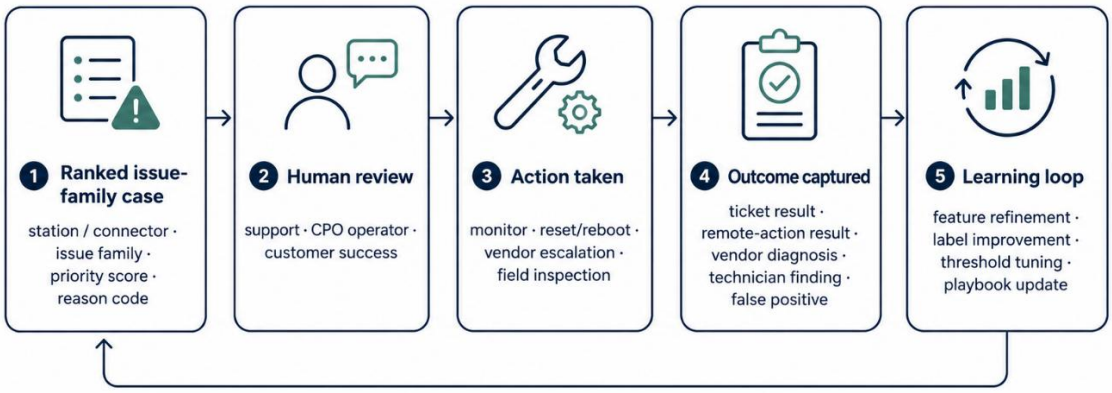

Figure 8.1 Feedback-loop design for maintenance-intelligence learning.

The key point of **Figure 8.1** is that maintenance intelligence should be treated as a learning system, not a one-off model output. The model can rank literally issue-family cases, but operational teams need to confirm whether those cases were meaningful, actionable, duplicated, resolved automatically, or escalated. This feedback is what allows GreenFlux to improve from “this connector should be reviewed” toward “this type of case usually requires this action under these conditions.”

This also means that ownership cannot sit only with the data or engineering team. Maintenance intelligence requires cross-functional ownership because the product value depends on the connection between data signals, user workflow, troubleshooting action, and operational outcome. **Table 8.3** summarizes the proposed ownership model.

The minimum feedback captured should be simple enough for operational use but structured enough for future model improvement. GreenFlux does not need a perfect feedback system from the first release, but it does need a consistent way to connect ranked cases with what happened afterwards. **Table 8.4** summarizes the minimum

feedback fields that would make the maintenance-intelligence loop useful for future machine-learning development.

Table 8.3 Ownership and feedback responsibilities.

Actor	Role in MVP	Feedback/data responsibility	Decision responsibility
Product owner/ product manager	Defines MVP scope, issue-family priority, and release gates	Collect user needs, product constraints, and adoption feedback	Decides which issue families enter MVP, pilot, or later-stage roadmap
Data/ Engineering team	Builds feature tables, issue-family flags, ranking logic, dashboards, and model monitoring	Tracks model performance, drift, data quality, feature reliability, and feedback-label availability	Decides whether a signal is technically usable for ranking, explanation, or model testing
Support team	Reviews whether ranked cases are operationally useful dan troubleshooting-friendly	Provides ticket outcomes, repeated-issue context, remote-action results, and escalation patterns	Validates whether reason codes and suggested routes fit real troubleshooting work
Customer success team	Connects the MVP to CPO workflow and adoption needs	Captures CPO feedback, workflow friction, usability issues, and training needs	Helps decide whether outputs are understandable and usable for customers
CPO operation teams	Uses dashboards and review queues in daily operations	Provides action taken, reviewed/not reviewed status, usefulness feedback, and operational outcomes	Confirms whether ranked cases help prioritize work and reduce troubleshooting uncertainty
Vendor/OEM or field-service partner	Provides technical interpretation for vendor-specific and field-level findings	Provides diagnosis, component replacement, firmware issue, installation issue, or no-fault-found feedback	Confirms whether issue-family interpretation aligns with technical reality
Leadership/ roadmap owner	Ensures business alignment and resource commitment	Review pilot outcomes, customer value, and risk exposure	Decides whether to scale from MVP to broader product capability or automated alerts

Table 8.4 Minimum feedback fields for future model improvement.

Feedback category	Example fields	Why it matters
Case review outcome	Reviewed, not reviewed, ignored, duplicate, escalated	Shows whether the ranking is useful in real workflow
Action taken	Monitor, remote reset/reboot, unlock connector, remote stop/start, support review, vendor escalation, field inspection	Links ranked issue families to operational responses
Remote-action result	Action attempted, success/failure, recovery time, repeated issue after action	Help distinguish remotely recoverable issues from escalation cases
Ticket linkage	Ticket ID, ticket category, issue owner, time opened, time closed	Connects model output with operational case management
Vendor diagnosis	OEM diagnostic category, firmware issue, known error, hardware issue, no issue confirmed	Improves vendor-aware issue-family mapping
Field-service finding	Component replaced, installation issue, cable issue, grid/site issue, no fault found, repair confirmed	Provides stronger validation for physical root-cause interpretation
False-positive label	Not actionable, repeated noise, contextual signal, resolved automatically, not reproducible	Reduces alert fatigue and improves threshold design
Resolution status	Resolved, unresolved, recurring, escalated, pending vendor, pending field visit	Helps evaluate whether review queues improve operational outcomes
Time-to-resolution	Time from ranked case to review, action, resolution, or closure	Measures operational value beyond model performance
Recurrence after action	Issue returned within 1 day, 3 days, 7 days, or 30 days	Help distinguish temporary recovery from persistent problem

These feedback fields should also be used as decision gates. GreenFlux should not move directly from dashboarding to automated predictive alerts. Each stage should be tested against operational usefulness, feedback quality, false-positive tolerance, and user trust. **Table 8.5** proposes decision gates for scaling the MVP.

Table 8.5 Decision gates for scaling from review queue to predictive alerts.

Gate	Question	Go/no-go implication
Gate 1: Issue-family clarity	Do users understand what each issue family means and where it appears?	If no, refine issue-family naming, documentation, and entity-level explanation before expanding

Gate	Question	Go/no-go implication
Gate 2: Dashboard usefulness	Do dashboards help users interpret raw logs faster than current views?	If yes, expand dashboard coverage to more issue families; if no, simplify views and labels
Gate 3: Review-queue usefulness	Do ranked cases help users decide what to review first?	If yes, move toward reason-code explanations; if no, improve ranking logic and workflow fit
Gate 4: Actionability	Do users know what next step to take after reviewing a ranked case?	If yes, add guided troubleshooting routs; if no, co-design playbooks with support and CPO operational teams
Gate 5: Feedback quality	Are outcomes captured consistently enough for model learning?	If yes, test calibration and thresholds; if no, improve ticket linkage and feedback capture first
Gate 6: False-positive tolerance	Are users still willing to trust and use the ranked output?	If yes, plot limited alerts; if no, keep human review and adjust thresholds
Gate 7: Automation readiness	Are precision, calibration, outcome labels, and user trust sufficient?	If yes, test narrow predictive alerts; if no, remain at explainable prioritization stage

The feedback-loop design makes the implementation path more realistic It allows GreenFlux to start with a useful human-in-the-loop product while building the evidence needed for future automation. In the short term, this supports dashboarding, review prioritization, and clearer troubleshooting. In the longer term, it can support calibrated alerts, better model monitoring, and more reliable predictive-maintenance functions. The condition is that every ranked case should become a learning opportunity, not just a one-time model output.

8.3 Expanded adoption-risk register

The adoption of explainable uptime intelligence depends not only on model performance, but also on trust, workflow fit, data quality, ownership, and feedback discipline. A technically valid ranking model may still fail if users do not understand why a case is ranked, if too many weak cases are surfaced, or if ranked cases are not connected to real operational outcomes. For this reason, adoption risk should be managed as part of the implementation design, not as an afterthought.

The main adoption challenge is that maintenance intelligence sits between product, data, support, customer success, CPO operations, vendors, and field-service partners. Each group sees a different part of the issue: data teams see signal patterns, support teams see tickets, CPOs see operational disruption, and vendors or technicians see technical diagnosis. The risk register below translates these adoption challenges into practical safeguards and ownership responsibilities. **Table 8.6** highlights analysis on adoption-risk register for explainable uptime intelligence.

Table 8.6 Expanded adoption-risk register for explainable uptime intelligence.

Adoption risk	How it appears in practice	Potential impact	Safeguard	Primary owner
False confidence in model output	Users interpret ranked issue-family cases as confirmed diagnosis or proof of physical failure	Wrong escalation, unnecessary field visits, reduced trust when diagnosis is not confirmed	Use model-quality labels, reason codes, and clear wording that outputs support review prioritization, not confirmed root-cause diagnosis	Product owner; data/ engineering
Alert fatigue and false positives	Too many weak or low-confidence cases are surfaced to users	Users ignore the queue, trust declines, and the feature becomes another noise source	Start with high-evidence issue families, use top-K review queues, monitor false positives, and only expand after pilot validation	Product owner; support; CPO pilot users
Workflow mismatch	Ranked cases do not fit how support teams or CPO operators actually triage problems	The tool is analytically interesting but operationally unused	Co-design review queues and action routes with support teams, customer success, and selected CPO users before scaling	Product owner; support; customer success
Vendor-specific behavior	A diagnostic pattern from one vendor is generalized to other vendors or charger types	Misleading issue-family mapping and reduced model transferability	Preserve vendor, model, AC/DC, station/connector, and protocol context in dashboards and model evaluation	Data/ engineering ; product owner
Weak feedback labels	Ranked cases are not linked to tickets, remote actions, vendor diagnosis, or field outcomes	Model cannot improve and cannot move toward calibrated alerts	Capture review outcome, action taken, remote-action result, ticket linkage, vendor diagnosis, technician finding, and false-positive labels	Support; data/engine ering; customer success
Station/ connector over-aggregation	Signals are collapsed into one generic	Loss of diagnostic meaning and wrong	Keep station-level and connector-level queues separates, with entity-specific	Product owner; data/ engineering

Adoption risk	How it appears in practice	Potential impact	Safeguard	Primary owner
	charger-health score	operational interpretation	explanations and issue-family context	
Data quality and telemetry gaps	Some issues families cannot be validated because telemetry, vendor diagnostics, or session fields are incomplete	Weak confidence in voltage, metering, no-current-flow, or power-transfer interpretation	Track data coverage, missingness, telemetry gaps, and vendor-field consistency before alerting	Data/ engineering
Over-automation too early	GreenFlux moves from ranking to automated alerts before calibration and trust are sufficient	Premature product claims, customer frustration, and reputational risk	Use decision gates, keep humans in the loop, and move to alerts only after feedback labels and false-positive tolerance are validated	Leadership; product owner
Unclear internal ownership	The model exists, but no team owns issue-family definitions, review workflow, feedback capture, or product gates	Fragmented implementation and slow adoption	Assign a maintenance-intelligence product owner and a cross-functional working group across product, data, support, and customer success	Leadership; product owner
Customer trust and explainability gap	CPO users do not understand why a station or connector is ranked	Low adoption and low willingness to act on ranked cases	Provide reason-code explanations, model-quality labels, entity-level context, and clear next-step guidance	Product owner; customer success
Ambiguous action routes	Users see a ranked case but do not know whether to monitor, reset, escalate, or inspect	Ranking does not translate into operational value	Add guided troubleshooting routes and validate them with support and CPO users	Product owner; support

Adoption risk	How it appears in practice	Potential impact	Safeguard	Primary owner
Uncontrolled expansion of issue families	Too many issue families are added before the first ones are validated	Product complexity increases and evidence quality becomes uneven	Start with a narrow pilot, then expand only after usefulness, feedback quality, and trust are demonstrated	Product owner; leadership
Model drift and changing charger behavior	Firmware updates, vendor changes, or new charger models alter signal behavior	Model performance degrades over time	Monitor performance by vendor, model, firmware, AC/DC type, and entity level; revalidate periodically	Data/ engineering
Misalignment with CPO value perception	The feature highlights technically interesting cases that CPOs do not consider operationally important	Low perceived value and weak commercial adoption	Track whether ranked cases reduce troubleshooting uncertainty, speed up triage, or improve escalation decisions	Customer success; product owner
Insufficient documentation and training	Users do not know how to interpret issue-family labels, reason codes, or model-quality levels	Misuse, confusion, and inconsistent adoption	Provide short user guidance, issue-family definitions, and examples of correct interpretation	Product owner; customer success

The most critical first-release risks are false confidence, alert fatigue, workflow mismatch, and weak feedback labels. These risks should be controlled before GreenFlux expands from dashboarding into broader review queues or automated alerts. A narrow release with strong issue-family boundaries is preferable to a broad release that creates noise or overstates diagnostic certainty. **Table 8.7** highlights risk-control priorities by implementation stage.

The adoption-risk logic supports the staged implementation path recommended in the main report. GreenFlux does not need to wait until every risk is eliminated before developing maintenance intelligence. Instead, it should start with a narrow, explainable, human-in-the-loop product and use each reviewed case as feedback for improving the system. The goal is to build trust gradually: first by making raw charger signals more

interpretable, then by helping users prioritize reviews, and only later by testing calibrated alerts where the evidence, feedback, and workflow fit are strong enough.

Table 8.7 Risk-control priorities by implementation stage.

Implementation stage	Main risks	Control focus
Issue-family dashboard pilot	Misinterpretation of issue families; over-aggregation	Define clear issue-family labels, separate station and connector context, and show model-quality boundaries
Review queue pilot	Alert fatigue; weak prioritization value	Start with high-evidence candidates, use top-K queues, and validate whether ranked cases help users decide what to review first
Reason code and action-route testing	Explainability gap; ambiguous next steps	Test whether users understand why a case is ranked and what action should follow
Feedback-loop validation	Weak feedback labels; missing outcome data	Link ranked cases to tickets, remote-action outcomes, vendor diagnosis, field findings, and false-positive labels
Limited alert testing	Over-automation; false positives, trust loss	Test alerts only for narrow issue families with sufficient calibration, outcome labels, and user tolerance
Product scaling	Vendor-specific drift; workflow mismatch across customers	Segment by vendor/model, AC/DC type, and CPO workflow; monitor performance and adoption continuously