# White Paper

## SAFEGRID
## Smart Fault Prediction
## 2.0

# From Detection to Decision:

Asset Health Intelligence using Safegrid Smart Fault Prediction 2.0

*How Safegrid Smart Fault Prediction 2.0 transforms partial discharge data into ranked, actionable maintenance intelligence*

## 1. Precursor Events and Predictive Maintenance

### Understanding Precursor Events

Power grid assets — conductors, joints, terminations, switchgear, and transformers — degrade over time due to a combination of electrical stress, thermal stress, mechanical vibration, moisture, and ageing insulation materials. Before a catastrophic fault occurs, assets rarely fail without warning. Instead, they exhibit a progressive pattern of measurable precursor events: subtle physical and electrical phenomena that signal deteriorating insulation integrity well in advance of failure.

Precursor events can take many forms and chief among these partial discharges (PDs) have emerged as one of the most reliable and sensitive leading indicators of emergent breakdown in medium- and high-voltage equipment.

### Low Energy Events (eg. Partial Discharges) as a Leading Indicator

Low energy events, including partial discharges (referred to herein as a "PD"), are a localised electrical spark that only partially bridges the insulation between conductors. PDs occur within cracks, contaminants, or loose interfaces inside the insulation system of a cable or piece of switchgear. Each discharge is extremely brief and releases a small, extremely fast burst of electromagnetic energy.

While an individual partial discharge event causes only minor insulation damage, the cumulative effect of repeated discharges erodes insulation steadily and accelerates degradation. Over months or years PD activity can intensify in both rate and energy until the insulation reaches breakdown, resulting in a critical fault. PD activity can precede a fault by days, months or even years — making it an important indicator for predictive maintenance programmes.

PDs are not all equally easy to detect. High-energy discharges may be captured by conventional sensors operating at lower sampling rates. However, the early-stage, low-energy PDs that are most informative from a condition-monitoring perspective require high-frequency sensing to resolve accurately. This distinction is central to the value Safegrid delivers.

### Predictive Maintenance: Positive Outcomes for Grid Operators

Traditional maintenance strategies for distribution networks rely either on time-based schedules (inspect every N years regardless of condition) or on reactive repair following a fault. Both approaches carry significant costs. Time-based maintenance can focus unnecessary work on healthy assets while potentially missing rapidly deteriorating ones. On the other hand, reactive maintenance is the most expensive and disruptive mode of operation - repairs don't happen until a failure occurs.

Predictive maintenance — the practice of monitoring condition indicators to intervene only when and where needed — changes this economics fundamentally. The positive outcomes are well-documented across the industry:

- **Improved SAIFI (System Average Interruption Frequency Index):** By identifying and remediating fault-prone assets before they fail, utilities reduce the number of sustained interruptions experienced by customers, directly improving this key reliability metric.
- **Lower Operations and Maintenance (O&M) costs:** Targeted interventions replace broad, scheduled sweeps or patrols. Maintenance crews are dispatched when there is genuine reason to act, reducing wasted callouts and unnecessary asset replacement.
- **Extended asset lifespan:** Catching and correcting problems at an early stage — a loose joint, degrading insulation, a developing gap — preserves asset integrity and extends serviceable life, deferring capital expenditure.
- **Reduced outage duration:** When faults are anticipated and assets pre-emptively taken out of service for repair, planned outages replace unplanned ones. Planned outages can be scheduled for low-demand periods with prepared crews, dramatically reducing customer impact.

**INSIGHT**

Studies have found that condition-based and predictive maintenance programmes **can reduce unplanned outages by as much as 75% and cut overall maintenance costs by 20-30%** compared with time-based and reactive maintenance regimes.(1) The key enabler is sensor acuity: the earlier a degradation trend can be detected, the wider the response window and the lower the intervention cost.

**Customer Case:** Enhancing Grid Visibility and Fault Detection

## 2. Challenges of Decision Making Based on Precursor Events

The promise of predictive maintenance based on PD monitoring is well established. The practical challenge, however, lies not in detecting precursor events but in deciding what to do about them. Utilities and network operators that deploy PD sensing infrastructure often find themselves facing a new and unexpected problem: an overwhelming volume of signals that is difficult to act upon.

### Alarm Fatigue and False Alarms

When monitoring systems detect and report every PD event without discrimination, the result is a continuous stream of alarms. Operators quickly learn that the vast majority of alerts do not correspond to imminent faults: they may indicate benign background noise, anomalous transient events, or low-level activity from assets that have been stable for years. The natural (and often rational) human response is to become desensitised to alerts.

Alarm fatigue is one of the most significant barriers to the effective use of predictive maintenance systems. When operators stop trusting their alerts, they stop acting on them. The monitoring system becomes a liability: it consumes resources without delivering value, and genuine warning signals may go unheeded in the noise.

### Increased O&M Cost for Low-Impact Callouts

Even where operators do respond to alarms, undifferentiated alert systems generate callouts that consume maintenance budget without proportionate benefit. A crew dispatched to investigate a low-level PD signal on a remote rural feeder may return with no actionable finding, having

[1] https://www.powermag.com/cut-downtime-and-costs-with-predictive-maintenance-pdm-heres-how/

incurred vehicle, labour, and administrative costs. Multiply this across dozens of alerts per month and the cost implications become material.

The challenge is compounded when alerts lack context about magnitude or urgency. Without knowing whether a PD trend is accelerating or stable, whether it represents an isolated event or a persistent pattern, and whether the affected asset serves critical load, maintenance planners cannot rationally prioritise. Everything looks equally urgent — or equally ignorable.

### Localised Problems vs System-Level Impact

A further challenge arises from the disconnect between where a problem is located and what its consequences would be if the asset were to fail. A developing fault in a densely loaded urban feeder may carry far greater impact than a similar defect on a branch serving a handful of rural customers: without impact context, both generate the same kind of alert.

Grid operators must balance the need to address local asset deterioration with a broader understanding of how each asset fits within the network topology and the downstream consequences of its failure. This is inherently a system-level question that individual alerts cannot answer in isolation. The most effective predictive maintenance regimes are therefore those that connect asset condition signals to network-level consequence modelling.

**Customer case:** Fast Detection Leads to Fast Action

## 3. How Safegrid Detects Precursor Events

### High-Frequency Sensing: Capturing What Others Miss

Safegrid's sensing platform achieves the necessary sensitivity for detecting low energy events by operating at sampling rates exceeding 1 MHz. This high measurement rate allows it to specifically capture the fast-rise-time, high-frequency current and voltage patterns characteristic of Partial Discharges (PDs). Many conventional grid monitoring systems operate at lower sampling rates (often between 1–10 kHz), which, while suited for generalised protection systems, are fundamentally unable to capture the subtle, short-duration signals associated with incipient insulation degradation, thus lacking the sensitivity needed for early detection. By sampling at over one million times per second, Safegrid sensors can capture subtle PD pulses. This difference in data collection sensitivity and measured frequency extends the observable window of the degradation lifecycle, providing operators with significantly more lead time to act, which represents the difference between a planned fix and an unplanned outage.

### Spatial, Temporal, and Energy-Level Pattern Recognition

Raw PD data on its own is not inherently informative. The analytical value lies in identifying patterns across three dimensions:

1. **Spatial patterns:** Where in the network are events occurring, and are they concentrated at specific locations or distributed across the system? Clustering of events at a particular cable span, joint, or termination is a key indicator of localised defects.
2. **Temporal patterns:** How is the event rate changing over time? A stable, low-level background is categorically different from a steadily or sharply increasing trend. Temporal pattern analysis reveals whether a problem is nascent and slow-moving, or active and accelerating.

3. **Energy level patterns:** What is the magnitude and distribution of individual discharge events, and how is this evolving? Increasing discharge energy, particularly when accompanied by an increasing rate, indicates advancing insulation deterioration.

Together, these three dimensions allow Safegrid to move beyond simply reporting that PD activity exists and toward characterising its nature, trajectory, and significance.

### Smart Fault Prediction 1.0: A Foundation for Visibility

Safegrid's first-generation Smart Fault Prediction platform (SFP 1.0) brought this multi-dimensional PD intelligence to market, making the location, timing, and waveforms of precursor events visible to network operators through a structured interface. For the first time, operators could see not just that faults had occurred, but where PD activity was accumulating and how patterns were developing over time.

SFP 1.0 represented a substantial advance over purely reactive monitoring. However, the question of whether a particular spatial, temporal, or energy pattern warranted immediate action — or whether it should be monitored further or deferred — remained a human judgement call. The platform provided the data; translating data into decisions remained the operator's responsibility.

We are pleased to announce that Safegrid has delivered a significant capability advance of this already-powerful toolkit by providing actionable guidance for prioritising attention to precursor events.

**Customer case:** Intermittent Earth Faults Detected in Challenging Weather

## 4. Smart Fault Prediction 2.0 (SFP 2.0): Prioritising Precursor Events

SFP 2.0 introduces an automated prioritisation engine that translates Safegrid's rich PD dataset into a ranked, actionable list of network locations requiring attention. Rather than presenting operators with raw signals and leaving interpretation to human intervention, SFP 2.0 automates the analytical work and presents only what matters — in the order that it matters.

*SFP2.0 Map View showing localised ongoing network events with priority ranking*

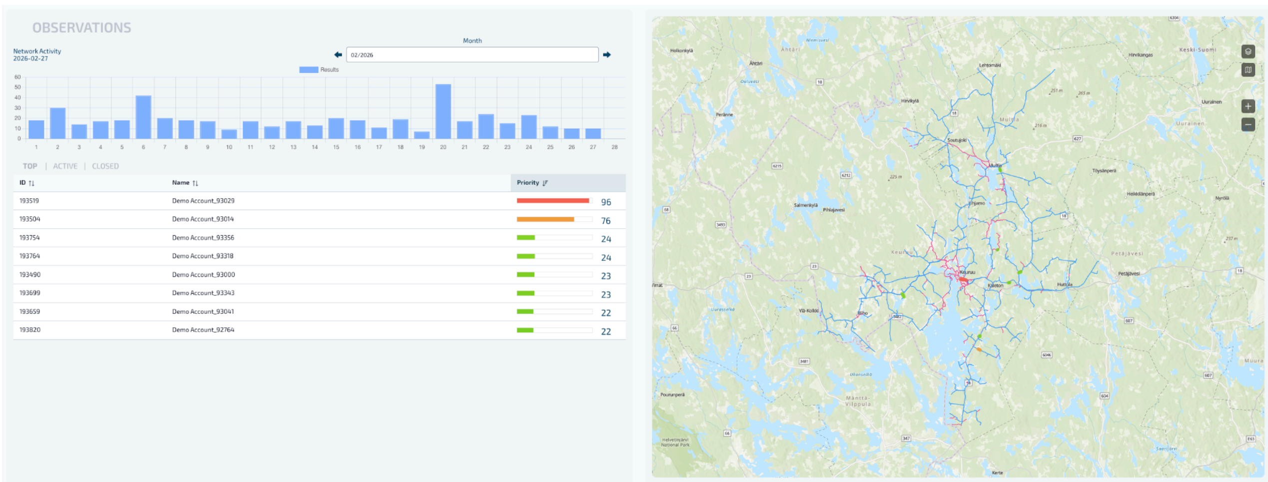

## Ranked Event States: Top, Active, Closed

Every identified event cluster in SFP 2.0 is assigned a dynamic state reflecting its current priority status. Locations are classified as ***Top*** (highest priority, requiring prompt attention), ***Active*** (ongoing PD activity that warrants monitoring and near-term planning), or ***Closed*** (activity has subsided or has been resolved). This simple, three-tier classification scheme gives operators an immediate, unambiguous view of where their attention should be directed.

The classification isn't static: it evolves continuously as new data arrives and patterns develop. A location may progress from Active to Top as activity intensifies, or transition to Closed as an asset stabilises. This dynamic ranking ensures that the priority list always reflects the current condition of the network, not an outdated snapshot from weeks or months ago.

*SFP2.0 Event View showing an event cluster on a single feeder*

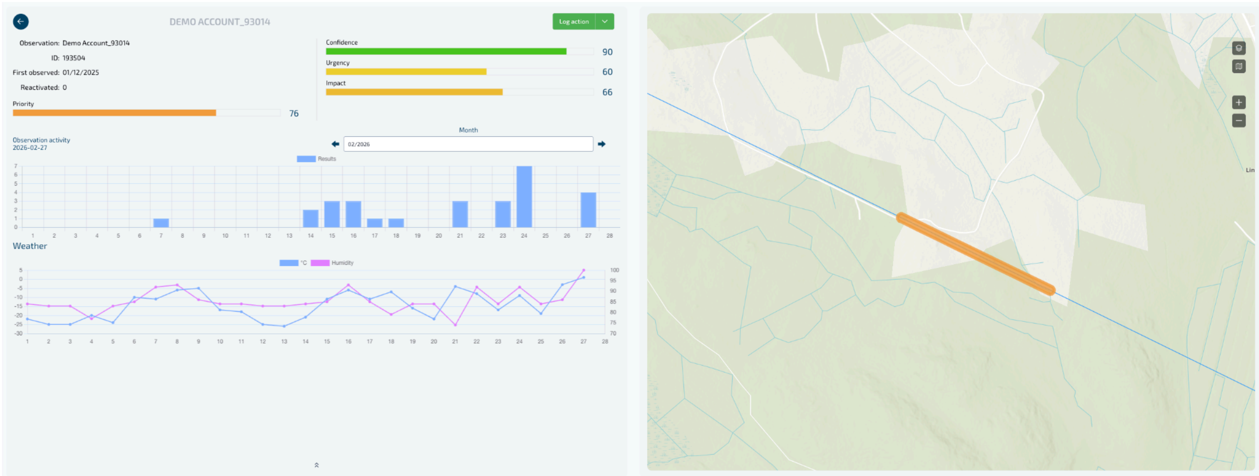


## The Three Dimensions of Priority: Confidence, Urgency, and Impact

The prioritisation engine in SFP 2.0 is built around three independently calculated dimensions that are combined to produce an overall priority score. Each dimension captures a distinct and important aspect of the situation.

1. **Confidence: Signal Quality and Statistical Reliability**

   Not all PD observations are equally credible. Electrical environments contain various sources of noise and interference that can produce signals resembling genuine partial discharges. SFP 2.0 applies rigorous signal quality assessment to evaluate the reliability of each observation. This includes statistical analysis of signal distributions: genuine localised insulation defects produce characteristic patterns of discharge activity, whereas background noise and interference tend to produce random, non-localised distributions that are statistically distinct.

   By incorporating confidence as an explicit dimension of the priority score, SFP 2.0 ensures that low-confidence observations do not generate spurious high-priority alerts. Only events that meet the platform's quality threshold are elevated for operator attention, directly addressing the false-positive problem that undermines alarm credibility in first-generation systems.

2. **Urgency: Rate of Change and Energy Trajectory**

   Urgency captures how rapidly a situation is developing. It is calculated from two components: the absolute rate of PD activity at a location, and the rate at which that activity rate is changing over time. A location experiencing a sudden step-change in activity is treated as more urgent than one with a similar absolute level that has been stable for weeks.

The urgency calculation also incorporates the energy trajectory of discharge events. Increasing discharge energy — particularly when combined with an increasing rate — is a strong indicator of accelerating deterioration. Both rate and energy are assessed on a moving-average basis, meaning that the current state of a location reflects recent history without being dominated by transient spikes or long-past episodes. This stochastic, recency-weighted approach captures the actual current threat trajectory rather than a simple cumulative count.

3. **Impact: Network Consequence Modelling**

   Impact quantifies the consequences that would follow if the asset at a flagged location were to fail. It is calculated from the network topology and the downstream power consumption of all load branches connected through the affected point. A failure at a high-load, low-redundancy node carries a substantially higher impact score than a failure at a peripheral spur with minimal connected load.

Impact scoring is what makes SFP 2.0 a critical concern management tool instead of a simple condition monitoring system. It ensures that maintenance resources are directed not simply toward the most active PD locations, but toward those where the combination of deteriorating condition and potential consequences justifies priority intervention. Critically, impact scoring also plays a powerful role in reducing false positives: a genuine low-impact location, even with active PD signals, appropriately receives a lower relative priority than a high-impact location with moderate but accelerating activity.

## The Priority Score and Intuitive Grid Topology Interface

Confidence, Urgency, and Impact are combined algorithmically to produce a single, composite priority score for each monitored location. This score drives the ranked list of events displayed in the SFP 2.0 interface, with Top events prominently surfaced for immediate operator attention.

The priority list is presented within an interactive grid topology map that shows the spatial relationship of flagged locations within the actual network structure. Operators can immediately see not just what requires attention, but where it sits in the network, what it is connected to, and how it relates to other active events.

This intuitive visualisation is a meaningful differentiator. Many current partial discharge and predictive maintenance solutions present data in tabular or chart-based formats that require operators to mentally construct a network-level picture from disaggregated outputs. SFP 2.0 does this work automatically. The result is a platform that requires significantly less domain expertise to operate effectively — the system handles the analytical interpretation, and the operator sees only what they need to act. This is not a superficial UX feature: it is a fundamental shift in where the cognitive work happens.

*SFP2.0 Event Detail View showing an event cluster correlated with ambient conditions*

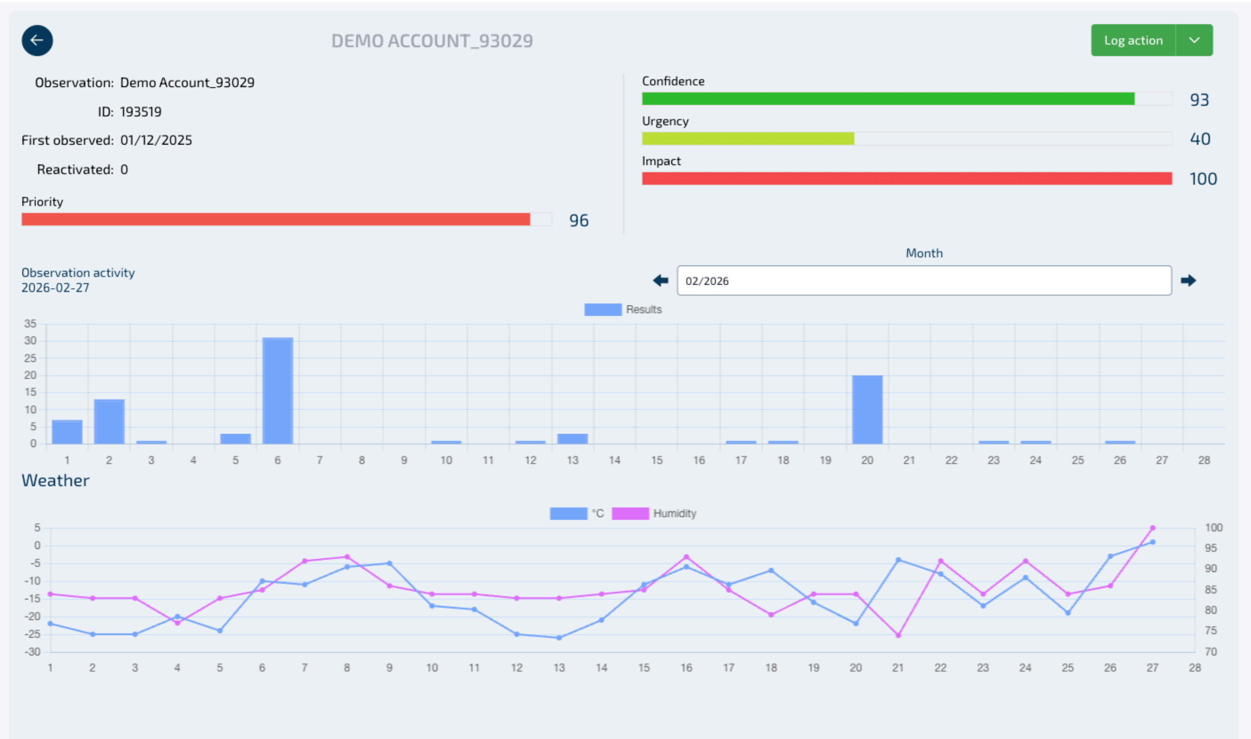


## Looking Ahead: Newer Capabilities in Smart Fault Prediction 2.0

The SFP 2.0 platform is designed as a continuously evolving system. Two significant capability extensions are in active development and will further enhance the platform's predictive and advisory functions.

- **Ambient Conditions and Cycle Forecasting for Proactive Scheduling**

  PD levels are not static: it varies primarily with ambient temperature and humidity, but also with load patterns and seasonal demand cycles. While load is not a definitive factor to Partial Discharge (PD), it can play a small role, as insulation that is performing adequately under normal conditions may reach a critical threshold during extreme temperatures. Conversely, a scheduled maintenance window that falls during a period of elevated grid stress may not be the optimal time to intervene.

  SFP 2.0 incorporates ambient condition and event forecasting to allow the platform to anticipate periods of peak aactivity and recommend pre-emptive action in advance of those periods. The forecast horizon is deliberately calibrated to be practically actionable: long enough to allow maintenance dispatch to be planned and resourced (more than a few hours), but short enough that the forecast remains reliable and relevant (less than a week). This window — broadly spanning one to five days ahead — sits at the intersection of operational utility and forecast confidence, avoiding the twin failure modes of being too short to act upon and too uncertain to trust. The result is a system that can prompt operators to get ahead of adverse conditions rather than responding to the consequences.

- **Asset-Level Fault Source Estimation**

  Knowing that a problem exists at a network location is valuable. Knowing what type of asset is most likely responsible is more valuable still, because it informs the nature of the maintenance response: the tools required for repair, the specialist skills needed, and the likely time required for repairs.

  SFP 2.0 will introduce asset-level fault source estimation, drawing on Safegrid's accumulated dataset from deployments across diverse network environments. Statistical models trained on this data provide probabilistic guidance on the most likely source of observed PD activity.

**Customer case:** Smart Fault Prediction Feature Helped Detect Potential Weak Spots

**Customer case:** Several Hundred Thousand Saved: Preventing Major Faults with Safegrid

## 5. Benefits and Business Case

### Operational and Financial Impact

The combined effect of SFP 2.0's Confidence, Urgency, and Impact framework — together with the intuitive prioritised interface — delivers tangible operational improvements across several dimensions. The table below on page 7 summarises modelled and indicative outcomes for a utility deploying SFP 2.0 compared with an unassisted or first-generation monitoring approach.

**INSIGHT**

Norwegian grid operator Griug had observed irregular signals via Safegrid monitoring, enabling a quick busbar fix that prevented a major outage and saved hundreds of thousands of NOK.

This proactive approach exemplifies how predictive maintenance cuts unplanned outages by up to 75% and O&M costs by 20-30%

### OPEX Impact: From Reactive Spend to Targeted Investment

The most direct financial benefit of SFP 2.0 is the reduction in wasted maintenance expenditure. In the absence of reliable prioritisation, maintenance teams respond to alerts that turn out to be low-consequence or non-actionable, and dispatch resources to locations that do not require intervention. The cost of these false positives accumulates rapidly: vehicle costs, labour costs, lost productivity for skilled technicians, and the administrative overhead of managing inconclusive callouts.

By filtering low-confidence events, scoring urgency on the basis of genuine trend acceleration, and weighting priorities by downstream impact, SFP 2.0 dramatically reduces the proportion of callouts that yield no actionable finding. Maintenance resources are concentrated on the events that matter most. For a utility managing hundreds of monitoring points across a medium-voltage distribution network, the reduction in unproductive callouts can represent a material OPEX saving in the first year of deployment.

The broader OPEX benefit comes from the shift from reactive to proactive maintenance posture. A fault that causes an unplanned outage triggers a cascade of costs: emergency crew dispatch (often at premium rates), extended outage duration, customer compensation, regulatory penalties, and reputational impact. Identifying and remediating the same developing fault weeks in advance, during a planned maintenance window with prepared materials, typically costs a fraction of the reactive equivalent. SFP 2.0's prioritisation engine is specifically designed to make this shift operationally reliable and repeatable.

### Reliability and Compliance Outcomes

SAIFI improvement — reducing the frequency of sustained customer interruptions — is a core regulatory metric for distribution network operators in most jurisdictions. Proactive fault prevention through predictive maintenance directly improves SAIFI performance. SFP 2.0 accelerates this improvement by ensuring that the predictive maintenance programme focuses on the highest-impact potential faults first, maximising the reliability benefit per maintenance Euro spent.

Asset lifespan extension is a secondary but significant benefit. Early-stage detection and remediation of developing faults preserves insulation integrity and prevents the progressive damage that eventually necessitates full asset replacement. Deferring capital asset replacement — even by a few years per asset — generates substantial capital expenditure savings at the portfolio level.

### A Platform That Grows With Your Network

Because SFP 2.0's prioritisation models are continuously updated with new data, the platform's accuracy and confidence improve over time as more observations are accumulated from each site. Asset-level fault source estimation becomes more precise as the platform builds familiarity with the specific characteristics of the local network. The forthcoming ambient and activity forecasting capabilities will add a proactive scheduling dimension that further compounds the value of the monitoring investment. SFP 2.0 is not a point-in-time product: it is an intelligence platform that compounds in value.

**Customer case:** Preventing Unnecessary Trips with Safegrid's Intelligent Grid System®

**Customer case:** Cable Fault Pinpointed by Safegrid's Intelligent Grid System®

## Taking Action

The question for power grid operators is no longer whether to monitor for precursor events. The evidence base for PD monitoring as a predictor of asset failure is established, and the regulatory and commercial pressures to improve reliability are intensifying. The question is whether the monitoring programme you operate translates into reliable, actionable decisions, or whether it generates data that is difficult to act upon and confidence that erodes over time.

Smart Fault Prediction 2.0 is built to answer this question definitively. It combines Safegrid's market-leading high-frequency sensing capability with an automated prioritisation engine that does the analytical work on your behalf, ranking events by confidence, urgency, and network impact. The results are presented in an intuitive, user-friendly interface that makes the right action immediately obvious.

We invite you to engage with us to understand how SFP 2.0 can help you realise the benefits of a more predictive maintenance regime. Whether you currently have a PD monitoring scheme in place or are evaluating predictive maintenance solutions for the first time, our team will work with you to show the tangible benefits that SFP 2.0 can provide.

See all our global customer reference stories from the following link:
https://safegrid.io/customer-references/

Disclaimer

**Safegrid Oy**
Otakaari 5, 02150 Espoo | Finland |
contact@safegrid.io
www.safegrid.io