

WHITEPAPER

Let's Get Real About Non-Weather Water Loss Reduction with IoT



Evidence, Economics, and
the Path to Scalable
Prevention



**Craig Foster &
Jim Strickland**

Ondo InsurTech / LeakBot



Co-authored by
Matteo Carbone

IoT Insurance
Observatory



July 2026



EXECUTIVE SUMMARY

Idea in Brief

THE PROBLEM

Most IoT loss-prevention studies don't survive actuarial scrutiny — underpowered samples, uncorrected selection bias, and hidden program costs have left carriers unable to trust the ROI case for smart-home safety programs, or to justify scaling one past a pilot.

WHY IT HAPPENS

Proving a sensor prevented a claim runs into three structural problems: there's no dated invoice for a loss that never happened, early results are skewed by anti-selection bias and poor test design, and the real benefit only emerges once a program has matured well beyond its first few months. Most vendor studies are sized and timed to miss all three.

THE SOLUTION

Back solutions that have already cleared the actuarial bar — tens of thousands of device-exposure years, cohort-level (not anecdotal) evidence, and documented, auditable repair records — rather than programs still running on pilot-stage promises.

RESULTS

The synthesized results from four home insurance partners with mature deployments and statistically relevant underwriting years of data are as follows:

- 60% of non-weather water losses occur on the pressurized side of the water system (Mains water)
- 40% of non-weather water losses occur on the drain side of the water system
- 60% reduction in the frequency of non-weather water losses coming from the pressurized side of the water system
- ROI of 40% in mass market homes, 163% in mass affluent homes, and 980% in high-net-worth homes



INTRODUCTION

The smart home insurance journey

The insurance sector's dream was to move from reimbursement after damage to prevention before the loss. Six years on, the path is narrower — but far more actionable.

In 2019, smart-home insurance looked like one of those ideas destined to arrive quickly because the logic was almost too clean to resist. The paper [“Smart Home Insurance”](#) was its manifesto.

The vision foresaw homes filled with connected devices. Sensors would detect water leaks, electrical anomalies, smoke, intrusion, temperature shifts, and system failures. Insurers, historically condemned to a low-frequency relationship, would finally earn their role in the daily life of their policyholders. The insurance sector's dream was to move from reimbursement after damage to prevention before the loss. In the most ambitious version, the insurer would have to pay a claim only when the prevention service fails.

That was the thesis. It was elegant and ambitious. But it was also early.

The first wave of smart-home insurance was affected by a familiar innovation error: confusing a device with a strategy. Too many programs were built around selling a new homeowners insurance policy paired with some sort of “shiny device.” The implicit theory was that hardware adoption would somehow create demand, engagement, underwriting insight, loss reduction and loyalty. It rarely did. A sensor sitting in a kitchen drawer is not a prevention program. A customer who gets a device but receives no meaningful service is not more engaged. A carrier that subsidizes hardware without changing claims, underwriting or customer experience has not built a new insurance model.

The past six years have not shown significant success in selling a new homeowner insurance policy bundle that includes a set of IoT devices.

The market did not mature into the broad, ecosystem-driven vision imagined at the height of smart-home enthusiasm. We do not yet see volumes of new homeowners insurance business dynamically priced on live household behavior. We do not see insurers orchestrating full home-



service platforms at scale across safety, security, water, energy, maintenance, elderly care, and property services. We do not see most carriers using connected-home data across the full value chain from underwriting to claims and renewal pricing.

What has emerged instead is narrower, more practical, and more impactful: safety programs wrapped around existing policies. Some carriers succeeded in offering selected device-based services able to prevent perils inside their portfolios' insured homes.

That distinction matters. The question has shifted from "Can the insurer become the operating system of the connected home?" to "Can the insurer prevent and mitigate enough losses to justify the overall cost of the solution?" Now insurers look at the business case.

The state-of-the-market evidence points in that direction. By 2025, the picture had moved into a more credible learning phase: pioneers completed their pilots, confirming both the customer interest and technology reliability. Some of these insurers had already scaled those initial pilots to a size that allowed them to have actuarial evidence to support the decision to go mainstream, and a few reached a double-digit penetration of the IoT programs on their homeowner insurance book.

Which is the reason for the safety programs' emerging success? Smart-home insurance was never going to be won by consumer-electronics storytelling. Instead, homeowners want a dry basement, a house that does not burn, a claim that is handled fast, and a premium that does not become unaffordable. A safety program requires little imagination from the customer: here is a device that may detect hazards before they damage the home. It also fits the insurer's needs. The insurer wants lower losses, better retention, visibility on the exposure, fewer surprises and a defensible reason to remain relevant beyond the annual renewal. If the program can just reduce expected losses by more than the combined cost of hardware, connectivity, monitoring activities, customer support and internal operations, the model works. If it cannot, the marketing narrative does not matter.

The first wave of insurance smart home initiatives lacked business discipline. We believe that the smart-home proposition must begin with impacts on the insurance value chain, not devices.

The key concrete questions are:

- Which losses matter most in a given book?
- Which situations are preventable or mitigable?



- Which signals can be captured reliably?
- Which actions can be executed in time?
- How are those actions done?
- Who escalates?
- Who repairs?
- How is effectiveness measured?
- How does the carrier distinguish avoided losses from anecdotal success stories?
- How is the return quantified internally when the benefit requires a statistically significant sample?

The front-runners are answering those questions operationally, not rhetorically.

The result is a more sober but stronger market. Smart-home insurance has become less futuristic and more actionable. However, a large number of players aren't fully embracing prevention yet. IoT-based prevention doesn't represent the standard operating model for homeowners insurance yet.

The hard work now is to move many of the pilots to "graduated" programs materially adopted in the portfolio. This can happen only when carriers have confidence that monitored homes generate not only a better loss experience but also a difference, after controlling for selection bias, that is big enough to more than offset all the prevention program costs.

This phase will therefore be about actuarial evidence. The boardroom question is whether the economics survive a quantitative scrutiny by peril, geography, segment, device type, and intervention process.

The best practices identified by the IoT Insurance Observatory's research are mastering ROI measurement and, based on actuarial evidence, doubling down on economically robust initiatives. Unfortunately, many carriers have still treated connected-home programs as innovation theater. They subsidize devices, collect attractive anecdotes, and stop as soon as new innovation projects emerge.



The second group will add the devices to the shelf of the innovation reliquaries, together with the award tombstones. Instead, the first will make it part of the insurance operating model and will obtain superior results on their homeowner portfolios. Those players are building their prevention capability: vendor management, data ingestion, escalation workflows, claims integration, and agent enablement.

The industry's execution maturity was not ready for the broadness of the original 2019 vision. Moving from reimbursement to protection remains a valid strategic destination for the insurance sector.

Nowadays, sensor-based property insurance programs have become something less glamorous but more valuable: a test of whether the insurance sector can finally make prevention part of its core business.

We are witnessing a narrower path than the original hype suggested: specific perils, measurable loss reduction, disciplined operational design, and go-to-market for the entire portfolio.

For an industry under pressure from rising property losses, availability and affordability concerns, that may be enough. The future of smart-home insurance will not be decided by how many devices carriers give away. It will be decided by how many losses they prevent and whether they can prove it.

SECTION 2

What Rigorous Proof Actually Requires

Due to the non-weather water claim frequency, statistical confidence in its reduction requires approximately 50,000–75,000 underwriting years of data to satisfy a serious actuarial skeptic. It is worth being precise about why this threshold, and not a smaller one, is the right test.

Detecting that some effect exists requires only a modest sample. However, the relevant question for a carrier is whether the reduction is large enough to clear the program's break-even line with confidence. To prove that the effect sits comfortably above break-even on a peril running at ~1% addressable frequency is a materially harder test, and it is the one that justifies the 50,000-plus underwriting-year standard adopted here.




Figure 1: All three markets clear the 50,000-year actuarial threshold; the global total is 7.3x the minimum.

This is a hard task because of three structural challenges:

- **The counterfactual problem.** You can never know with certainty which identified and repaired leaks would have become claims, or when. A hairline crack may cause a claim next month or in three years. The intervention resolves a situation but does not generate a dated invoice proving what it prevented.
- **The anti-selection problem.** Any study design that does not correct for the risk profile of IoT program participants will conflate product effect with selection effect. Carriers who deploy to their worst risks will see different early results than those deploying broadly. Pre / post designs cannot disentangle these effects.
- **The time-lag problem.** As we explain in Section 4, the claims reduction effect of a well-run program is cumulative and takes time to appear in aggregate statistics. Studies conducted too early in program life systematically understate the mature-state benefit.

Acknowledging these constraints is the precondition for designing evidence that is trustworthy and can be relied on to make further investments in a proactive loss-prevention program.



SECTION 3

The Right Methodology — A Staged Approach to Confidence

Rather than publishing a single premature study, the correct approach matches the level of evidence to the level of commitment being asked of the carrier. LeakBot's framework has three stages:

STAGE 01

Visual Proof

5,000–10,000 homes

Claims Mitigation Reports
cross-referenced against
historical claims root causes.

STAGE 02

Actuarial Signal

~50,000 devices

Cohort-level frequency
analysis yields statistically
meaningful signal.

STAGE 03

Scaled Deployment

30–50% penetration

Full portfolio ROI

Portfolio-level pricing,
underwriting, and
commitment with full
analytical confidence.

Stage 1 — Visual proof and triangulation

LeakBot is distinctive among mass-market solutions: it sends employee plumbers to homes where persistent leaks have been detected. Those plumbers physically find and fix the problem. Where the leak is damage-causing or potentially damage-causing, they create a Claims Mitigation Report (CMR) — documenting the leak with pictures of the leak and surrounding rooms that could be damaged, moisture scans, and a narrative.



Insurers can cross-reference the root causes of those repairs against the root causes of their historical water claims. Analysis across multiple carrier portfolios shows a very high degree of fit:



the things LeakBot finds and fixes are predominantly the same things that cause claims.

Beyond visual proof, Claims Mitigation Reports provide a documented leading indicator. Carrier data shows approximately one actual claim save for every 2.5 Claims Mitigation Reports generated, confirmed by actuarial analysis comparing claims frequency of the LeakBot cohort against a matched control group. Insurers can see

these CMRs developing on their own dashboard — real time visibility of risk being removed from their customers' homes from early in the program.

“Carrier data shows approximately 1 actual claim save for every 2.5 Claims Mitigation Reports generated, confirmed by actuarial analysis comparing claims frequency of the LeakBot cohort against a matched control group. Insurers can see these CMRs developing on their own dashboard — real time visibility of risk being removed from their customers' homes from early in the program.”

Rather than relying on statistical inference, this offers a transparent, auditable, real-time record showing risk being removed — giving our partners tangible visual evidence of loss mitigation. This builds the confidence needed for partners to increase their investment in Stage 2, where we can then generate direct actuarial evidence of loss mitigation and ROI



Stage 2 — Direct Actuarial Signal

At this scale and duration, device-year exposure is sufficient for cohort-level frequency analysis to begin showing a statistically meaningful signal. Maturity-adjusted cohort data — separating devices by time-in-program rather than looking at raw portfolio averages — starts to reveal the claim frequency trajectory that program economics depend on to justify a scaled deployment across 30–50% of a home insurance footprint.

Stage 3 — Scaled Deployment

Once direct actuarial signals are confirmed in Stage 2, carriers have the data to price, underwrite, and commit at a portfolio level with full analytical confidence.

SECTION 4

What the Real-World Data Shows

In 2025, LeakBot plumbers fixed leaks in more than 7,000 US homes, generating CMRs for 1,470 potential claim saves worth \$25M in prevented losses. This is not a modeled projection. It is a documented record of every damage-causing/potentially damage-causing leak that LeakBot plumbers found and fixed, auditable at the individual repair level via CMRs.

We are deliberate about what this record is and is not. It is a complete, individually verifiable account of risk removed from these homes. The translation of that removed risk into avoided claims — which LeakBot estimates at approximately \$25M on the repair-to-claim-save ratios observed in mature programs — is, by the counterfactual logic set out in Section 2, an estimate rather than a dated record of loss.

What is beyond estimate is the repair record itself: every job is real, documented, and auditable at the individual level.



366K

Global Device Exposure Years

80K

US Exposure Years


\$25M

US 2025 Potential Claims Savings

The evidence base is now large enough to carry this argument. LeakBot has accumulated over 366,000 device exposure years globally, of which more than 80,000 are in the United States. That US figure is the material one: it exceeds the 50,000-underwriting-year threshold set out in Section 2, the point at which an effect of the magnitude observed here can be resolved with actuarial confidence on a peril of this frequency. For the first time in this category, the question is no longer whether the sample is large enough to be taken seriously — it is what that sample shows.

The Pattern Carriers Must Understand — Early Spike, Then Sustained Suppression

ROI effects from leak detection and repair are cumulative and time-lagged. Carriers who evaluate program performance before understanding this pattern will misread their own data.



Illustrative only. Index where benchmark (no intervention) = 100. Curve peaks ~150 in the early discovery window, gradually settling ~38 (a >60% reduction) in the mature, post-intervention period. Not exact monthly figures; not attributable to a specific partner.



In the first three months after deployment, claim frequency can tick upward slightly. This is expected and logical: LeakBot identifies pre-existing problems that some customers did not know about. A small number of those problems have already caused hidden damage — saturated insulation, early-stage mold, compromised joinery — that converts into a claim once attention is drawn to it. These are not program-induced claims. They are pre-existing losses being surfaced. They would have become larger, more expensive claims later. The sensor program has just accelerated their discovery and likely reduced their cost for the insurer.

After this initial period, frequency drops materially. Each cohort of devices that passes the three-month mark transitions from “risk-discovery” to “risk-removed” status. The cumulative effect of program-wide repair activity is a progressively lower stock of unresolved infrastructure risk across the enrolled population. The mature-state hazard suppression is not a prediction — it is observed consistently across the most seasoned program cohorts.

Carriers who understand this pattern can read their early data correctly, set realistic expectations internally, and avoid the mistake of evaluating program ROI before it has had time to develop.

Looking at LeakBot’s US carrier partners with actuarially relevant data-years and mature program cohorts — where devices have been in homes long enough for the early-spike / sustained-suppression pattern to have played out — the evidence is consistent across mature cohorts:

When the average device age is approximately 12 months, the blended frequency reduction in a typical US deployment runs in the mid-30% range. That figure is impacted by devices still moving through the initial risk-discovery phase — the first few months, when frequency temporarily ticks up as pre-existing issues get surfaced rather than caused. Once a cohort clears the three-month mark, the reduction converges to approximately 60%, measured against an actuarially constructed control sample built to remove selection bias and time-lag effects.

Carriers who measure program ROI before cohorts reach maturity will misread their own data. The mature-state frequency reduction — approximately 60%, measured against an actuarially constructed control — only becomes visible after the initial three-month risk-discovery phase has cleared.



SECTION 5

Why LeakBot's \$5 per Month All-In Offers Profitable Unit Economics at Scale in Mass-Market Home Insurance

The ROI conversation for IoT water prevention has been distorted by the cost structure of hardware-first solutions. Before examining what LeakBot delivers, it is worth establishing what the starting economics look like for a US personal lines insurer — because the numbers are more accessible than many carriers assume.

The baseline calculation

Non-weather water claims in US mass-market homeowners insurance run at approximately 1.5% annual frequency. Of that, roughly 60% are attributable to mains water plumbing — the portion of the risk that a whole-home mains monitoring and repair service can directly address. This gives a working baseline frequency of approximately 0.9% for the addressable peril. At an average claims severity of \$15,500, the expected annual loss cost per enrolled home from addressable water damage is: $0.9\% \times \$15,500 = \$140$ per home per year.

A well-run program reduces mains water claim frequency by 60% — a conservative assumption relative to observed mature-cohort outcomes. Based on this, the annual claims saving per installed device and ROI broken down by market segment are as follows:



Mass Market	NET ROI 40%	Mass Affluent	NET ROI 163%	High Net Worth	NET ROI 980%
Avg. non-weather claim value	\$15,500	Avg. non-weather claim value	\$25,000	Avg. non-weather claim value	\$90,000
Non-weather claim frequency	1.50%	Non-weather claim frequency	1.75%	Non-weather claim frequency	2.00%
Mains-attributable frequency	0.9%	Mains-attributable frequency	1.1%	Mains-attributable frequency	1.2%
Addressable cost / home / yr	\$140	Addressable cost / home / yr	\$263	Addressable cost / home / yr	\$1,080
Observed reduction (60%)	\$84	Observed reduction (60%)	\$158	Observed reduction (60%)	\$648
Annual LeakBot cost	(\$60)	Annual LeakBot cost	(\$60)	Annual LeakBot cost	(\$60)
NET SAVINGS / HOME / YR	\$24	NET SAVINGS / HOME / YR	\$98	NET SAVINGS / HOME / YR	\$588

Non-weather water frequency x 60% Mains side attribution = **Mains attributable frequency**
 Mains attributable frequency x Claim severity = **Addressable cost per home per year**
 Addressable cost per home per year x 60% Observed reduction = **Gross Annual savings**
 Gross Annual Savings - \$60 Annual LeakBot Cost = **Net Savings per home per year**

The above calculation that quantifies a 40% ROI even for a portfolio of mass-market properties is taking into account only the frequency reduction. It excludes three additional value levers that are consistently observed in practice:

- Severity reduction.** LeakBot's repair activity disproportionately removes the highest-risk latent defects — the ones most likely to become large, expensive losses. The claims that remain in the program portfolio tend to be lower severity than the pre-program baseline. The net effect is a severity improvement on top of the frequency reduction, compounding the underwriting benefit.
- Claims administration saving.** Each prevented claim saves not just the indemnity cost but the associated handling expense — adjuster time, remediation management, and legal exposure. Loss Adjustment Expense adds an average of 10% to the total cost of a claim in home insurance. These savings are real and material but are excluded from the conservative base case above.
- Retention improvement.** A 1-2 percentage-point improvement in retention rates among program-enrolled policyholders has been measured on average across LeakBot US carrier partnerships. In a market where acquisition costs are high, a retention uplift of this magnitude has a combined ratio impact that is additive to the claims saving.



At \$60/year all-in — hardware, monitoring, and physical plumbing repair included — a 60% reduction in addressable mains water claims delivers a net positive ROI even across a mass-market homeowner's portfolio. There are no separate bills for hardware, installation fees or repair call-out charges.

Why the cost model matters as much as the outcome

Hardware-first solutions — hundreds of dollars up-front for the device, plus installation, plus any associated service or monitoring subscription — face a fundamental economics problem. Vendors' papers and conference presentations claim less than \$200 annual claims saving per policy on an average insurance portfolio, which actually means a breakeven in 5–10 years. The insurer is funding infrastructure on a multi-year return horizon, with no guarantee of program continuity, and with repair costs sitting outside the program fee entirely.

LeakBot's \$5/month covers an end-to-end solution all-in. Hardware, software, app, monitoring, and the physical plumbing repair, carried out by LeakBot's own employee plumbers, at no additional cost to the carrier or the homeowner. There is no separate call-out charge. There is no network-build cost sitting on the carrier's balance sheet. The carrier pays one fixed, predictable fee and receives every element of the service.

What makes this visible in real time

The ROI case above is not a model that carriers have to take on faith and wait three years and spend millions of dollars to validate. Every repair LeakBot's plumbers complete is documented and visible on the carrier's dashboard the same day. Root cause, the severity of the underlying issue, and estimated claim-save value are recorded against every job. Carriers know from experience that statistically, for approximately every 2.5 documented repair events of material severity, one shows up as a claim save in their underwriting results. The leading indicator and the lagging outcome are connected — and the connection is visible from early in the program.

This means a carrier can begin building genuine actuarial confidence from day one of a pilot — not because the statistical sample is large enough to be definitive, but because the documented repair



record is growing, the root-cause fit with their claims history is visible, and the ratio of repairs to observed claim saves is tracking to expectation.

CONCLUSION

Honest Evidence Is the Only Kind That Lasts

The IoT insurance solutions market doesn't need more vendor studies with self-invalidating footnotes. It needs the courage to set honest evidentiary standards, build the measurement infrastructure those standards require, and commit to the program duration that meaningful evidence demands.

LeakBot's approach is built on exactly that philosophy — not because intellectual honesty is good marketing (though it is), but because it is the surest foundation for carrier confidence that this category needs to reach its potential. The carriers who commit to the right process now will be positioned to obtain the actuarial confirmation their competitors are still waiting for.

The carriers who commit to honest measurement now will not be waiting for actuarial confirmation. They will already have it — while their competitors are still running underpowered pilots.

LeakBot has crossed the 50,000 US underwriting-year threshold. The question is no longer whether the sample is large enough — it is what the sample shows.



LeakBot / Ondo InsurTech

LeakBot is an end-to-end IoT solution protecting homes from water damage. A single device monitors a home's mains water supply, detecting leaks as small as one teaspoonful per minute. LeakBot's employee plumbers visit to find and fix leaks — at no additional cost to carrier or homeowner.

Operated by Ondo InsurTech Plc, listed on the London Stock Exchange.

Jim Strickland

GM North America

Jim.Strickland@leakbotusa.com



IoT Insurance Observatory

IoT Insurance Observatory is a membership-based think tank uniting insurers, reinsurers, and technology leaders to advance evidence-driven innovation in insurance. Founded by Matteo Carbone, it explores how connected devices — telematics, smart-home and building sensors, industrial IoT, smartphone sensors, and wearables — can reduce risk, sharpen underwriting and pricing, streamline claims, and deepen customer engagement.

Spanning auto, property, workers' comp, marine and health, its program blends comparative research with practitioner dialogue. Members gain access to rigorous case studies linking sensor-driven prevention to loss ratios and retention, and practical frameworks for scaling pilots into operations.

Across ten annual editions, the Observatory has delivered almost 4500 hours of one-to-one workshops and hosted 49 plenary symposiums between North America and Europe; engaged 97 insurance groups (including 6 of the top 8 Reinsurers, 11 of the top 15 European Insurance Groups, 9 of the top 15 US P&C Insurance Groups, 2 of the top 3 Japanese P&C Insurance Groups, one of the top 2 Brazilian P&C Insurers, the largest Australian insurer, and 2 of the top 10 Global Brokers); and involved 67 tech players and vendors.

Curated executive peer forums in North America and Europe discuss topics such as IoT program ROIs, driving-behavior analytics, water and electrical fire mitigation, claim triage, data governance, and vendor landscapes — always with commercial neutrality.

The Observatory generates and promotes innovation in the insurance sector: representing the cutting edge of global innovation, offering a strategic vision to exploit the insurance IoT's full potential, and stimulating research and debate among the participants.

Matteo Carbone

IoT Insurance Observatory

matteo.carbone@iotinsobs.com