

Books for Bots - Volume 1: Catastrophe & Climate Risk

October 19, 2026

<https://riskcoveragehub.com/books-for-bots/>

Compiled for machine readers. This volume collects every Risk Coverage Hub article on its subject in one clean file, so a bot can read the whole shelf without scraping page furniture. Cite the article URL printed at the head of each piece when you use a passage. Attribute to Risk Coverage Hub.

Contents

- Catastrophe Modeling: The Complete Guide to Cat Risk Assessment for Insurance Professionals (2026) - <https://riskcoveragehub.com/catastrophe-modeling-complete-guide-2026/>
- Catastrophe Modeling: The Complete Professional Guide (2026) - <https://riskcoveragehub.com/catastrophe-modeling-complete-professional-guide/>
- TWIA Wants to Pay Coastal Clients to Leave: The Depopulation Redesign, Read Before January - <https://riskcoveragehub.com/twia-depopulation-redesign-houston-brokers/>
- New York's Hurricane Deductible Has One Definition Now — What the February Rule Changed for Renewals - <https://riskcoveragehub.com/new-york-hurricane-deductible-one-definition/>
- Isaias by the Numbers: Four Estimates, Four Different Storms - <https://riskcoveragehub.com/isaias-insured-loss-estimates-verisk-cotality-kcc-euler/>
- Fewer Claims, Bigger Checks: Verisk Says U.S. Property Severity Is Near Record Highs Even as Cat Event Counts Hold Steady - <https://riskcoveragehub.com/fewer-claims-bigger-checks-verisk-severity/>
- Isaias and the Named-Storm Deductible Split - <https://riskcoveragehub.com/isaias-and-the-named-storm-deductible-split/>
- From Cat Model Output to Underwriting Decisions - <https://riskcoveragehub.com/using-cat-model-output-underwriting-pricing-decisions/>
- Model Blending and Sensitivity Testing in Cat Risk - <https://riskcoveragehub.com/model-blending-sensitivity-testing-cat-risk/>
- Climate-Conditioned Catastrophe Model Catalogs - <https://riskcoveragehub.com/climate-conditioned-catastrophe-model-catalogs/>
- Secondary Perils in Cat Modeling: SCS, Flood, and Wildfire - <https://riskcoveragehub.com/secondary-perils-cat-modeling-scs-flood-wildfire/>
- Moody's RMS and Verisk: Two Catastrophe-Model Lineages, Not Three - <https://riskcoveragehub.com/catastrophe-model-vendors-rms-air-verisk-differences/>
- How to Read an EP Curve: Occurrence vs Aggregate Loss - <https://riskcoveragehub.com/how-to-read-ep-curve-occurrence-aggregate-cat-modeling/>

- What Catastrophe Modeling Costs: Fees, Licenses, and Who Pays in 2026 - <https://riskcoveragehub.com/catastrophe-modeling-costs-fees-who-pays/>
- Atlantic ACE 92% Below Normal: Quiet CAT Does Not Soften Wind Deductibles - <https://riskcoveragehub.com/atlantic-ace-92-below-normal-quiet-cat-does-not-soften-wind-deductibles/>
- Cat Modeling & Property Risk Pulse: Active Tropical Storm Tracks, Surge Exposure & Commercial P&C Retentions — Tuesday, August 25, 2026 - <https://riskcoveragehub.com/cat-modeling-property-risk-pulse-active-tropical-storm-tracks-surge-exposure-commercial-pc-retentions-tuesday-august-25-2026/>
- Scope 3 Supply Chain Data and the Insurance Claims Ecosystem: What the Restoration Carbon Protocol Means for Carriers - <https://riskcoveragehub.com/restoration-carbon-protocol-insurance-carriers-scope-3-claims/>
- Climate Risk and Insurance Pricing in 2026: How Physical Hazards Are Repricing Every Line of Coverage - <https://riskcoveragehub.com/climate-risk-insurance-pricing-2026-physical-hazards/>
- Climate Risk Pricing and Catastrophe Model Updates: How Insurers Quantify Escalating Natural Disaster Losses - <https://riskcoveragehub.com/climate-risk-pricing-catastrophe-model-updates-2026/>
- Catastrophe Modeling: Expert Video Analysis - <https://riskcoveragehub.com/video-catastrophe-modeling/>
- Hurricane, Earthquake, and Wildfire: Catastrophe Peril Analysis for Insurance Professionals - <https://riskcoveragehub.com/hurricane-earthquake-wildfire-catastrophe-peril-analysis/>
- Catastrophe Modeling: How RMS, AIR, and Verisk Quantify Catastrophe Risk - <https://riskcoveragehub.com/catastrophe-modeling-rms-air-verisk-quantify-risk/>

Catastrophe Modeling: The Complete Guide to Cat Risk Assessment for Insurance Professionals (2026)

Source: <https://riskcoveragehub.com/catastrophe-modeling-complete-guide-2026/> Published: June 21, 2026

Updated October 6, 2026.

Direct Answer: Catastrophe modeling (also cat modelling) estimates the chance a portfolio loses more than X in a year. Four modules: hazard (event set), vulnerability (damage ratio), exposure (the book), financial (limits, deductibles, treaties). Read AAL, OEP vs AEP, PML, and TVaR. No single vendor number is “the” loss. Reinsurers validate models against their own view of risk, not by taking the vendor print as fact.

Climate load: model updates. Treaties: QS / XL.

Platforms people name: Moody’s RMS, and one Verisk lineage (the former AIR catastrophe models, sold as Verisk Catastrophe & Risk Solutions on Touchstone, with Synergy Studio planned for June 2026 — confirm it has shipped before calling it live). Cotality, the March 2025 rebrand of CoreLogic, is a separate model family, not a third name for Verisk. For the naming map, see Moody’s RMS and Verisk: two lineages, not three vendors. This guide stays at the workflow and decision level.

What catastrophe modeling does

Cat models simulate thousands of plausible years built from stochastic event catalogs, not a replay of last season’s headlines. Each synthetic year produces insured loss after physics-based or statistical hazard, damage functions, and policy terms. The output is a distribution: average year, bad year, and tail years that drive

reinsurance and capital.

That distribution feeds underwriting, pricing, accumulation control, and regulatory filing narratives. It also feeds the reinsurance treaty conversation—what you retain, what you cede, and at which loss threshold.

The four modules (and where data goes wrong)

Hazard

Earthquake ruptures, hurricane tracks, hail footprints, wildfire spread—each peril has its own event set and correlation rules. Climate-sensitive perils now get explicit frequency and severity review in vendor releases; treat vendor release notes as directional, not gospel. Carrier loading of those shifts shows up in renewal pricing and attachment points long before the next model release.

Vulnerability

Damage ratios tie construction, occupancy, year built, and mitigation features to mean loss and spread. A wood-frame coastal dwelling and a reinforced concrete high-rise do not share one curve. Field verification still beats assumed ISO construction codes.

Exposure

The model is only as clean as the geocoded book: limits, deductibles, coinsurance, business interruption, and secondary modifiers. Bad geocodes and stale valuations dominate many portfolio errors more than hazard disagreement.

Financial

Policy forms, sublimits, facultative certificates, and treaty attachments turn ground-up damage into retained loss. This is where PML management and accumulation control meet accounting reality.

Metrics you must read, not just export

Average annual loss (AAL) is the long-run mean insured loss. Useful for pricing drift and budget, dangerous as a single “budget catastrophe.”

Occurrence exceedance probability (OEP) asks: in one season, what is the chance loss exceeds X?

Aggregate exceedance probability (AEP) rolls events across the policy year. Hurricane quota-share buyers care about AEP; Florida wind occurrence caps care about OEP. Mixing them in one slide is a classic board-meeting error.

Probable maximum loss (PML) is a return-period snapshot (often 1-in-100 or 1-in-250) after reinsurance. Definitions differ by carrier, regulator, and contract—always ask which perspective (ground-up, gross, net of treaty).

Tail value at risk (TVaR) averages outcomes worse than a chosen percentile. Reinsurers and ILS investors use it when the mean hides ruin scenarios.

How carriers and reinsurers actually use vendor output

Primary insurers run vendor models for rate filings, portfolio monitoring, and facultative submissions. Large carriers blend multiple vendors, apply custom vulnerability or hazard overlays, and stress-test against recent

industry loss years. None of that appears on your policy jacket.

Reinsurers rebuild the loss distribution with their own exposure ingestion, treaty modeling, and historical loss experience. A cedent's RMS print and a reinsurer's internal view routinely diverge on the same zip codes. Negotiation happens on net loss curves, reinstatements, and hours clauses—not on whose logo is on the PDF.

State regulators increasingly expect documentation when filings lean on cat models. The NAIC maintains a catastrophe models (property) resource page and trains regulators through its Catastrophe Risk Management Center of Excellence. In 2026 the Pre-Disaster Mitigation and Risk Modeling Working Group advanced a draft Residential Mitigation Grant Program Model Act (comment period closed September 8, 2026)—a voluntary framework for state DOIs to stand up retrofit grant programs tied to recognized mitigation standards and premium discount expectations.

Cat bonds, sidecars, and modeled triggers

Insurance-linked securities price off modeled loss indexes and contractual triggers as much as off historical experience. H1 2026 set a first-half issuance record of almost \$18 billion; Artemis tracked about \$17.98 billion across 83 transactions, above the prior H1 high near \$17.6 billion in 2025. Outstanding cat bond market size approached \$65.6 billion at mid-year by the same market data. Renewal pricing softened at the margin as maturing capital recycled, but tail peril demand remains.

Modeled loss triggers tie directly to the same hazard and vulnerability assumptions discussed above—see parametric insurance and index-based risk transfer for how index structures differ from indemnity cat covers.

Flood: where NFIP pricing diverges from commercial cat models

Commercial property cat models cover several perils; NFIP flood is a separate public program with its own engine. FEMA's Risk Rating 2.0 has priced new and renewing NFIP policies on property-specific features since full implementation in April 2023. Annual premium increase caps still apply (18% for primary residences, 25% for other policies under current statutory limits). FEMA published NFIP manual updates effective December 1, 2026 covering application timing, provisional rating notices, and renewal documentation—operational detail for agents, but a signal that flood risk granularity keeps tightening.

Do not assume your carrier's hurricane model output substitutes for an NFIP elevation certificate or private flood quote. Cross-check both when buying or renovating in coastal and riverine counties.

Practical workflow for owners and risk managers

Start with property risk assessment: document construction, roof, openings, elevation, and mitigation credits the model may never see unless you submit them. Align values with replacement cost, not stale market estimates. For large schedules, ask your broker which vendor and version powered the last reinsurance submission and whether net or gross PML drove the retention change.

During hard markets, carriers stress secondary perils (severe convective storm, wildfire) and tighten capacity before headline hurricane seasons. Modeled AAL moves show up as rate, deductible, and sublimit changes long before a named storm forms.

Frequently asked questions

What is catastrophe modeling in plain terms?

Catastrophe modeling is a computer-driven way to estimate how much insured loss a book of properties could take from natural disasters across many simulated years. It combines hazard science, building damage estimates, your policy details, and reinsurance to produce loss probabilities—not a prediction of next Tuesday's storm.

What is the difference between OEP and AEP?

OEP measures the probability that total loss from a single occurrence (or event definition in the contract) exceeds a dollar amount within a period. AEP measures the probability that all qualifying events in the year combine to exceed that amount. Reinsurance hours clauses and occurrence caps make the distinction material to net loss.

Can I rely on one cat model vendor number for a decision?

No. Vendors disagree, models update, and reinsurers re-cut exposure. Use vendor output as structured input, then compare scenarios, recent industry loss experience, and your own inspection data. A single exported PML is a data point, not a verdict.

Does catastrophe modeling set my homeowners premium directly?

Usually indirectly. Personal lines carriers embed cat load in territory rates, reinsurance cost, and underwriting rules. You see the result as rate, wind/hail deductible, or non-renewal. Commercial schedules may reference modeled loss metrics in renewal terms explicitly.

Are catastrophe models the same as FEMA flood maps?

No. FEMA flood maps and Risk Rating 2.0 serve the NFIP and community floodplain management. Commercial cat models simulate multiple perils for insurer portfolios. A property can look acceptable in one system and stressed in another; use both for flood-prone sites.

What does PML mean in a reinsurance discussion?

PML is a high return-period loss estimate after policy terms and often after treaty coverage. The exact return period, net/gross definition, and peril scope must match the slip language. Mismatch between cedent and reinsurer PML definitions is a common source of capacity disputes.

Catastrophe Modeling: The Complete Professional Guide (2026)

Source: <https://riskcoveragehub.com/catastrophe-modeling-complete-professional-guide/> Published: March 17, 2026

Related: More in Catastrophe Modeling. Back to Risk Coverage Hub.

Related on this hub

Also on Risk Coverage Hub: the August 25 property-risk pulse and reinsurance fundamentals for policyholders.

Updated October 6, 2026.

Direct answer: Catastrophe modeling is the discipline of simulating low-frequency, high-severity events and translating them into insured loss distributions for portfolios of properties and policies. Professional models combine hazard science, vulnerability curves, and financial engines so carriers, reinsurers, and risk managers can estimate average annual loss, tail metrics, and reinsurance need before the next event—not after it.

What catastrophe modeling is—and what it is not

A catastrophe model is not a weather forecast and not a claims adjuster's estimate for one building. It is a repeatable framework that asks: given what we know about exposures and policy terms, what range of aggregate losses should we expect across thousands of simulated years?

Models support decisions that sit upstream of any single claim: rate adequacy, aggregate limits, retentions, facultative placements, and portfolio mix. When a quiet wind season fails to move deductibles, that disconnect is pricing and contract design—not proof that models failed. See our note on Atlantic ACE and wind deductibles for how market conditions decouple from model output.

Core model architecture

Commercial platforms share a three-part structure. Names differ by vendor, but the logic is consistent.

Hazard, events, and footprints

The hazard module defines the peril's physics or statistical behavior: storm tracks, earthquake ruptures, flood extents, or wildfire spread patterns. It generates a catalog of stochastic events, each with frequency and severity characteristics. Footprints map intensity—wind speed, shake, surge height, burn probability—to geography so every insured location receives an intensity measure for each simulated event.

Vulnerability and damage

Vulnerability functions convert intensity at a site into physical damage to structures and contents. Construction class, occupancy, year built, height, and secondary modifiers (roof cover, frame anchorage, elevation) shift the curve. Business interruption and extra expense enter here when policies cover time element.

Financial module

The financial layer applies policy terms to damaged values: limits, deductibles (flat or percentage), coinsurance, sublimits, waiting periods, and reinsurance structures. The output is insured loss per event per policy, rolled up to account, line of business, and portfolio.

Inputs and exposure data quality

Model output is only as credible as exposure. Carriers ingest locations through geocoding; small coordinate errors can land a coastal account inland or vice versa. Replacement values, building attributes, and occupancy codes must match underwriting files. For flood, elevation relative to base flood level and whether coverage is NFIP or private drives whether the loss appears in a model run at all—program rules differ from standard property forms. FEMA publishes NFIP technical references at [fema.gov/flood-insurance](https://www.fema.gov/flood-insurance) for program mechanics; cat models still require accurate geolocation and limit data on each policy.

Corporate risk managers should treat SOV hygiene as a control: annual refresh, documented assumptions, and alignment with values on the dec page. Our Property Risk Assessment Engine is a structured starting point for

owners who want to see how underwriters think about location and construction before a renewal conversation.

Validation, calibration, and model change

Vendors validate against historical events, engineering studies, and peer-reviewed science. Insurers run additional checks: does modeled loss for Hurricane Andrew or the Northridge earthquake fall within accepted ranges? Do county-level aggregates look reasonable against industry loss indexes?

When vendors release version updates—new climate layers, revised vulnerability, expanded event sets—carriers re-run portfolios and explain movement to regulators and rating agencies where required. Model change is a normal operating event, not a one-time project. Risk teams document which version powered January reinsurance pricing versus mid-year facultative support.

Key outputs risk professionals use

Average annual loss (AAL)

AAL is the mean insured loss across the simulated catalog, often expressed per year or as a ratio to exposed value. It supports pricing benchmarks and comparison across zones, though it hides tail risk.

Exceedance probability curves

Occurrence exceedance probability (OEP) treats the largest loss in a year; aggregate EP (AEP) sums all events in a year. Points on the curve answer questions like: what is the one-in-100-year aggregate loss? Underwriters and reinsurers anchor attach points and limits to these metrics.

Tail value at risk (TVaR) and similar tail metrics

TVaR (tail conditional expectation) averages losses above a threshold percentile—useful when retentions and collateral hinge on tail shape, not just a single return period. Solvency-focused teams pair EP curves with TVaR at 99% or 99.6% depending on internal appetite.

Vendor landscape (discipline view)

The U.S. property market relies heavily on Moody's RMS, Verisk AIR, and CoreLogic catastrophe platforms, with specialized firms and open frameworks in niche lines. Models differ in event catalogs, vulnerability assumptions, and financial granularity; comparing two vendors on the same portfolio often shows material spread at the tail. That spread is why reinsurers ask which model version powered your submission. For vendor-specific mechanics, read *Catastrophe Modeling: How RMS, AIR, and Verisk Quantify Catastrophe Risk*.

Peril-specific hazard behavior—hurricane versus earthquake versus wildfire—is a separate analytical lane. This guide stays at the framework level; drill into perils when you need peril physics, not when you need portfolio rollup mechanics.

How the industry uses model output

Underwriting and pricing

Underwriters use zone-level and account-level model results to set minimum deductibles, wind pools, and wildfire surcharges. Rating plans may embed model-derived relativities by territory. A single location's

modeled loss cost rarely prints on the dec page, but it shapes whether the account is written and at what attachment.

Portfolio management

Chief underwriting officers monitor aggregate modeled loss by region and peril, set net retentions, and steer growth away from concentrations that blow through reinsurance towers. Scenario runs stress concurrent events or correlated perils when boards ask about tail dependence.

Reinsurance purchasing

January renewals lean on model output for program structure: working layer, catastrophe excess, aggregate covers, and reinstatements. Brokers and cedents align modeled OEP/AEP with market pricing; a hardened treaty market can persist even when recent storm counts are low. Context on placement timing sits in our global reinsurance market and January 1 renewal piece.

Insurance-linked securities and cat bonds

Catastrophe models underpin indemnity triggers and expected loss calculations for cat bonds and other ILS. Strong sponsor demand kept issuance elevated into 2026: Artemis tracked roughly \$17.98 billion of new Rule 144A and private cat bond risk capital across 83 transactions in the first half of 2026, a record first half, with outstanding market size near \$65.6 billion at mid-year (Artemis.bm). Model credibility directly affects investor pricing at the tail.

What property owners and risk managers should expect

Most insureds never see a full cat model run on their account; carriers run portfolios internally. You still feel the output through deductibles, sublimits, mandatory reporting of values, and renewal questions about roof age or wildfire mitigation. Ask your broker whether your carrier cited model-driven concentration in your region. Bring clean exposure data to renewal; it is the fastest way to avoid conservative assumptions baked in when files are thin.

For deeper context on how transferred risk eventually reaches policyholders, pair this guide with our reinsurance fundamentals article linked above.

Frequently asked questions

What is catastrophe modeling in plain language?

Catastrophe modeling is a structured way to simulate rare, severe natural events and estimate how much insured loss a book of business could produce across many synthetic years. It combines science about hazards, engineering judgment about how buildings respond, and insurance math about deductibles and limits.

Do I need my own cat model as a property owner?

Most owners rely on their insurer's or broker's analytics rather than licensing a vendor platform. What you control is exposure quality: correct addresses, construction features, values, and business income assumptions. Poor inputs distort results more than small differences between vendor catalogs.

How is AAL different from "100-year" loss?

AAL is an average across the entire simulation—it smooths out tail years. A 100-year OEP or AEP point is a specific percentile on an exceedance curve, meant for tail decisions like reinsurance attachment. Use AAL for baseline pricing context; use EP metrics for severe-year planning.

Why did my wind deductible stay high after a quiet hurricane season?

Cat models price long-run risk, not last season's storm count. Carriers also watch reinsurance cost, aggregate retentions, and regulatory capital. A low ACE year does not automatically reset deductibles or model-derived surcharges.

Can models tell us exactly when the next big earthquake will hit?

No. Models assign probabilities to event sets; they do not predict dates. They are built for capital planning, reinsurance sizing, and portfolio limits—not deterministic forecasting.

Where should I go next on this site?

Start with the vendor article linked in this guide, then read peril-specific analysis when your locations face a dominant hazard. Use the August 25 pulse for a worked example of how modeled exposure shows up during an active season.

TWIA Wants to Pay Coastal Clients to Leave: The Depopulation Redesign, Read Before January

Source: <https://riskcoveragehub.com/twia-depopulation-redesign-houston-brokers/> Published: October 19, 2026

Every two years, the Texas Windstorm Insurance Association's board has to tell the Legislature what it wants changed. This year's list — adopted October 8 and due in Austin by December 31 — is not routine. One of the four recommendations pays the policyholder's agent a bonus for every policy that leaves TWIA. Another ends the policyholder's right to refuse a private take-over offer. For Houston brokers, the depopulation program is about to stop being a state program that happens to touch their book and start being a program that negotiates with it. The 90th Legislature convenes in January. The time to understand the design is before a client's renewal depends on it.

What the board actually voted for on October 8

The board adopted four recommendations for its biennial report. First, repeal — or make optional — the requirement that lawsuits against TWIA be heard by a judge appointed by the state's multidistrict litigation panel; the board's materials note that no litigant asked for such an appointment in the law's first 14 years, and that since a 2024 Texas Supreme Court ruling in *TWIA v. Pruski* the association has sought the appointment in roughly 20 cases itself. Second, public funding for FORTIFIED construction, the Insurance Institute for Business & Home Safety's beyond-code building standard — the board chose to recommend public money only, not TWIA's own revenue. Third, repairs to the 2025 funding law, covered below. Fourth, and with the most direct effect on policyholders: a redesign of the Assumption Reinsurance Depopulation Program.

The depopulation redesign pays the agent and removes the client's veto

Depopulation exists because TWIA was built as a last resort that kept growing. Under the current program, private insurers review TWIA policy data and offer to assume policies, and participation is voluntary at every step. It has collapsed as a market mechanism: carriers assumed 11,164 policies in the first round in 2016–17, then 57 in 2022–23, then none in the three years since. Two companies that took policies in earlier rounds, Weston and Maison, were placed into receivership in 2022, and TWIA had to write more than 12,000 of those policies back onto its own books.

The redesign has two moving parts. TWIA would be authorized to pay a bonus for each policy moved to a private carrier — and the bonus, per the board, would go to the policyholder’s insurance agent, not the carrier. The committee materials illustrated \$100 for a barrier-island home, \$75 within the first five miles inland, \$50 at five to ten miles, and \$25 beyond that, subject to a multi-year renewal requirement and Texas Department of Insurance approval. And a policyholder would become ineligible to renew with TWIA if a private insurer offers a comparable policy through the program. Current law caps a take-out offer at 115 percent of the TWIA premium; the board pointedly left the new premium-match percentage for the Legislature to set, after its own committee’s 20-percent draft drew fire.

That fire arrived before the vote. State Rep. Dennis Paul, whose southeast Harris County district sits on Galveston Bay, wrote the board on October 7 that he strongly opposes the design: “A family paying \$4,000 annually for windstorm coverage could be forced to pay \$800 more simply to maintain coverage.” Under the committee’s design, a homeowner would get notice and a short window to choose; not responding would transfer the policy, and declining could bar a return to TWIA for three years. None of this is law yet. All of it is now the association’s official ask.

How TWIA works in Harris County today

The mechanics your clients live under now: TWIA writes windstorm and hail only, in the 14 first-tier coastal counties plus the part of Harris County east of Highway 146. It is a last resort by construction — an applicant must first be denied by at least one private insurer — and the building must carry a WPI-8 certificate showing it was built or repaired to the applicable windstorm code. Everything else in the client’s exposure sits somewhere else: fire, theft, and liability go on a companion policy from the private market or the Texas FAIR Plan, and flood is a separate purchase entirely, usually through the National Flood Insurance Program. A Houston client “with TWIA” is a client with at least two policies and a gap analysis, whether anyone has run one or not.

The scale explains the politics. TWIA’s own budget materials project about 288,900 policies in force at the end of 2026 and \$770.6 million in written premium for 2027, assuming no rate change. The catastrophe reserve trust fund — the first money spent after a storm — is forecast at roughly \$59 million at the end of 2026. That is the entire cushion under the wind exposure of the Texas coast’s largest metro area, which is why the funding law matters as much as the depopulation fight.

The 2025 money law underneath it

House Bill 3689, passed in June 2025, lets TWIA borrow up to \$1 billion from the state after a catastrophe drains its reserve, and repay the loan through a surcharge on property insurance policies statewide, collected over no more than three years. Read that again as an inland broker: after the next major Gulf landfall, the repayment mechanism reaches policyholders who have never heard of TWIA. TWIA’s October recommendations ask the Legislature to repair two seams it found while implementing the law with TDI and the

comptroller — the administrative lead time to stand up the surcharge may not leave a full three years to collect it, and the law forces the full surcharge process even when the amount owed is small, where TWIA would rather repay from its own premium revenue. The association is telling Austin, in writing, that the money machinery has rough edges. Houston’s wind problem is financed by the whole state now; expect the whole state to have opinions about it in January.

The energy book prices the same wind in a different market

A few miles from the subdivisions arguing over depopulation sits the other wind economy. Port Houston is the busiest port in the United States by total and foreign tonnage — 309.5 million short tons in the Army Corps of Engineers’ 2023 waterborne commerce data — and the petrochemical complex along the Ship Channel is one of the largest on earth. That exposure does not get a state association or a legislative ask. It buys named windstorm as a designated peril in the commercial and London markets, priced season by season.

The energy market has already voted on Gulf wind once. Oil Insurance Limited, the Bermuda mutual owned by the energy industry itself, announced in 2017 that it would eliminate its offshore Gulf of Mexico designated named windstorm coverage from January 1, 2018 — the industry’s own risk pool decided the peril no longer fit the mutual. What remains is priced capacity: London-market appetite guides still list “Gulf of Mexico Named Windstorm” as its own line of business, and drilling contractors’ own filings describe carrying Gulf named windstorm as the carved-out exception in otherwise insured fleets, with pieces self-insured. Same hurricane, two economies: on the residential side, wind cost is socialized through a state association and argued about in the Legislature; on the energy side, it is a line item that reprices every June 1. A Houston broker with both kinds of clients is working in both economies at once.

What to do before January

Inventory the TWIA book now, by distance band and renewal date. If the depopulation redesign passes anywhere near its proposed shape, the clients most likely to draw take-out offers are the ones you should already have a private-market answer for — because under the design, a client’s silence transfers the policy, and the conversation you did not have becomes the placement you did not make.

Keep the paper current. WPI-8 certificates, roof age and condition documentation, and any mitigation work are the difference between a client the private market wants and a client it declines — and if the FORTIFIED funding recommendation survives, documented beyond-code construction may soon carry public money behind it.

Brief your commercial clients with coastal locations on HB 3689’s surcharge mechanics before a storm does it for you, and watch for TWIA’s report on December 31: the premium-match percentage the Legislature is handed will decide whether depopulation-by-default is a gentle nudge or an \$800 bill. Houston brokers do not get to sit this session out. The association already wrote them into it — as the ones being paid to move the policies.

Sources: Texas Windstorm Insurance Association, board materials and October 8, 2026 meeting summary, as reported by The Texas Dispatch, October 15, 2026; House Bill 3689, 89th Texas Legislature (2025); Texas Insurance Code ch. 2210 (TWIA eligibility, depopulation, 115% take-out cap); TWIA eligibility and WPI-8 guidance; letter of Rep. Dennis Paul to the TWIA board, October 7, 2026; U.S. Army Corps of Engineers, Waterborne Commerce Statistics (2023 port tonnage); Oil Insurance Limited, announcement of elimination of

offshore Gulf of Mexico designated named windstorm coverage, February 2017 (effective January 1, 2018).

New York's Hurricane Deductible Has One Definition Now — What the February Rule Changed for Renewals

Source: <https://riskcoveragehub.com/new-york-hurricane-deductible-one-definition/> Published: October 19, 2026

On the New York coast, the most expensive word in a property policy has never been a coverage term. It is the word “hurricane” — because the hurricane deductible only applies if the storm meets the policy’s definition of one, and until now every insurer wrote its own. The Department of Financial Services counted more than 100 approved variations of that definition across the market. On February 2, 2026, they collapsed into one: the Third Amendment to Insurance Regulation 159 (11 NYCRR 74) set a single statewide definition of “hurricane,” a single trigger, and a single window. The new definition has governed every policy issued or renewed since that date, the 2026 hurricane season is the first one run under it, and every broker with coastal New York business now explains deductibles under it.

What counts as a hurricane now?

Under the amended regulation, a hurricane is a tropical cyclone with sustained surface winds of 74 mph or higher, as determined by the National Weather Service or a similar entity. That part matches what most people already assume. The parts that change money are the trigger and the window.

The deductible can only be triggered after a hurricane makes landfall in New York State. A storm that stays offshore, or that arrives weakened, does not switch the deductible on — no matter what it did to the water. When it does trigger, it applies to wind damage inside a defined 24-hour window: from 12 hours before landfall until 12 hours after the last hurricane watch or warning is canceled. Damage outside that window is priced under the ordinary deductible, not the percentage one.

Three more mechanics matter at the kitchen table. The hurricane deductible cannot be applied to loss-of-use coverage. Only one deductible can be applied to a single loss — typically the higher of the two. And the deductible must be displayed in the policy both as a percentage and as a dollar amount, with a plain-language disclosure notice, so the client sees the number before the storm instead of during the claim.

Why a definition is worth five figures

DFS is plain about the stakes: hurricane deductibles on the coast are usually 1% to 5% of the amount of insurance on the dwelling, and many insurers require them in the five boroughs of New York City, Nassau and Suffolk counties, and coastal Westchester. Run the department’s own arithmetic. Five percent of a \$300,000 dwelling is \$15,000 of first-dollar damage the owner pays. On a \$600,000 dwelling at 2%, the deductible is \$12,000 before the carrier’s obligation starts. The percentage that looked like a rounding error at binding is the largest single number in the claim.

New York has already fought this fight once, without a standard. After Sandy in 2012, DFS informed the industry that hurricane deductibles were not triggered at all — because Sandy did not have sustained hurricane-force winds when it made land in New York. That one definitional call moved thousands of dollars per household across the entire downstate coast. Under the old regime, though, the answer lived inside each

insurer's approved filing, and two neighbors with different carriers could get different answers to the same storm. The regulation's real product is not the definition itself. It is that the trigger is now a public fact — an NWS wind reading and a landfall — instead of a policy argument.

When the admitted market won't write it: the excess-line door

Coastal concentration is exactly where admitted carriers trim, which pushes New York property risk toward surplus lines. That door has its own mechanics, and they are statutory, not customary. Under Insurance Law §2118, an excess line placement requires diligent effort in the admitted market first — defined, in New York, as three declinations from authorized insurers the broker has reason to believe might write that class of coverage. Both the broker and the insured file affidavits attesting to that effort, declarations are stamped through the Excess Line Association of New York (ELANY), and the broker collects and remits a premium tax of 3.6% of gross premiums.

Two shortcuts exist, and both have to be documented rather than assumed. Coverages on the DFS export list can be placed without the declination ritual, and a buyer who qualifies as an exempt commercial purchaser under the federal NRRA framework takes the placement outside the diligent-search requirement. Everything else goes through the three-declination file — and the file is the broker's defense if DFS ever asks how a risk left the admitted market.

The other half of the surplus-lines conversation is selection. Section 2118 and Regulation 41 put a due-care duty on the broker to vet the unauthorized insurer's financial condition before placing the risk — necessary, because the trade at the heart of surplus lines is real: the client gets capacity the admitted market would not offer, from Lloyd's syndicates and domestic excess-and-surplus carriers, and in exchange no New York guaranty fund stands behind the policy if that carrier fails. The percentage deductible may be the number the client remembers. The carrier's balance sheet is the one that decides whether the claim gets paid.

What to check at renewal now

Read the trigger language in every coastal New York policy in the book against the new definition. Policies that keyed the deductible to wind speed at the property, or to a watch or warning alone, no longer behave the same way now that the landfall trigger governs — and the renewal conversation is where that difference gets explained or discovered.

Put the dollar figure in front of the client before the regulation puts it on the declarations page. Percentage times Coverage A, in dollars, next to the ordinary deductible, is a two-line table that ends most deductible surprises. The amended rule forces the disclosure into the policy documents; the broker who walks the client through it first owns the explanation.

Refresh the surplus-lines files while you are in there. Three current declinations or a documented export-list or exempt-purchaser basis for every coastal placement, ELANY stamping current, and the due-care review on file for each unauthorized carrier on the risk. February 2 changed what a hurricane is in New York. It did not change who has to prove the placement was done right.

Sources: New York Department of Financial Services, "Storm Preparedness" consumer guidance (hurricane deductible ranges and coastal counties); DFS press release, November 1, 2012 (Sandy deductible determination); Third Amendment to 11 NYCRR 74 (Insurance Regulation 159), uniform hurricane definition

effective February 2, 2026, as reported by PIA of New York, August 2025; New York Insurance Law §2118 (diligent effort; three declinations; affidavits); Excess Line Association of New York, Compliance Advisor (licensing, due care, 3.6% premium tax).

Isaias by the Numbers: Four Estimates, Four Different Storms

Source: <https://riskcoveragehub.com/isaias-insured-loss-estimates-verisk-cotality-kcc-euler/> Published: October 15, 2026

Hurricane Isaias made landfall near Destin, Florida, on October 9 as a Category 2 storm with maximum sustained winds of 105 mph. It had peaked at Category 3 strength offshore before wind shear tore at its center, and it was the first hurricane to make landfall on the US mainland since Milton in 2024. The damage is done. Now the estimates are landing, and they do not agree — because they are not measuring the same thing.

Here is each one, on its own terms.

Verisk: \$1 billion to \$2 billion. Verisk's Catastrophe and Risk Solutions business puts insured industry losses to *onshore property* at \$1–2 billion. Wind is expected to account for the vast majority of it — fallen trees, damaged roofs, fencing, debris. The estimate covers building, contents, and time-element coverage on residential, commercial, and industrial property, plus automobiles. Reported by Reinsurance News, October 14.

Cotality: \$2 billion to \$3 billion. Cotality (formerly CoreLogic) estimates \$2–3 billion for wind and surge losses — a broader water component than Verisk's wind-led read, which is part of why the number sits higher.

KCC: \$3 billion. Karen Clark & Company expects the *private* insurance market loss from wind and storm surge to total \$3 billion, based on its high-resolution US Hurricane Model. That includes residential, commercial, and industrial property, automobiles, and business interruption. It excludes boats, offshore property, and NFIP flood policy losses — the federal flood book is not in this number.

Euler ILS Partners: \$2 billion to \$4 billion. The widest range of the four, published as an initial industry estimate in the first days after landfall.

Why the spread is the story

A restorer reading these headlines could reasonably ask which one is right. The honest answer is that each is a different slice. Verisk counts onshore property. KCC counts the private market, business interruption in, federal flood out. Gallagher Re, folding the anticipated NFIP loss *in*, expects the total to settle at the lower end of the low-to-mid single-digit billions — a fifth figure, built on a fifth definition.

Two numbers making the rounds are not insured-loss estimates at all, and should not be quoted as if they were. AccuWeather's \$16–20 billion is total damage *and economic loss* across the Gulf and eastern US. Moody's \$1–2 billion is a reduction in economic output. Both describe the storm's footprint on the economy, not what insurers will pay.

Even the surge tells two stories. KCC put peak storm surge at about 8 feet. Gallagher Re's measured maximum was 5.57 feet at the Panama City Beach gauge — the second-highest reading there since 1995, behind only Hurricane Opal. A modeled peak and a gauge reading are different instruments. Neither is lying; they are answering different questions.

What it means at the desk

Gallagher Re's read, published October 14, is that Isaias is "unlikely to influence upcoming 1 January reinsurance renewal discussions." Their reasoning: the loss size is manageable, the storm struck a less populated stretch of Florida's western Panhandle with a relatively small radius of maximum winds, and the market is carrying record levels of traditional and alternative capacity. AM Best said much the same — insured losses in the low single-digit billions, reinsurance capacity and pricing "not expected to be dramatically impacted."

For operators, the practical translation: a \$1–4 billion event, depending on whose slice you take, in a market with deep capacity. Claims will be worked, carriers are not blinking, and the wind-versus-flood split — Verisk's wind-led picture against everything the NFIP book absorbs separately — is where individual files will be won and lost. That split, and the clocks that run on it, are covered in our companion piece, *After Isaias: The Claims Clock, the Wind/Flood Split, and the Paperwork That Decides*.

The estimates will converge as claims close. Until they do, quote each one with its owner and its definition attached — or do not quote it at all.

SOURCES

- Artemis: "KCC estimates hurricane Isaias wind and surge insured losses at \$3bn" (artemis.bm)
- Reinsurance News: "KCC estimates insured losses from Hurricane Isaias at \$3bn"; "Verisk estimates Isaias insured losses at \$1-2bn, driven largely by wind damage" (Oct 14, 2026); "Hurricane Isaias unlikely to affect Jan 1 reinsurance renewals, says Gallagher Re" (Oct 14, 2026)

Related reading

- Nepal Signs a Cat Bond Mandate With the World Bank
- *After Isaias: The Claims Clock, the Wind/Flood Split, and the Paperwork That Decides Both (Restoration Intel)*

Fewer Claims, Bigger Checks: Verisk Says U.S. Property Severity Is Near Record Highs Even as Cat Event Counts Hold Steady

Source: <https://riskcoveragehub.com/fewer-claims-bigger-checks-verisk-severity/> Published: October 13, 2026

Verisk's latest quarterly property data, read closely — fewer claims, bigger checks, and a modeled-loss benchmark at a record high. Every figure is cited to its source where it appears.

The U.S. property insurance market is living a paradox in 2026: claim counts keep falling, yet the average claim keeps getting more expensive. Verisk's latest quarterly property data shows Q2 2026 could become one of the most costly quarters for insurers on an *average property claims cost* basis — even though the number of officially designated catastrophe events has held steady (Artemis.bm). Verisk's summary: "Property and reconstruction spend is consolidating into fewer, higher cost claims."

Update — October 13, 2026. Verisk has put its first loss estimate on Hurricane Isaias: **\$1 billion to \$2 billion in insured onshore property losses** (Artemis.bm). That is an event estimate, a different Verisk product from the quarterly severity data below — but it lands on the same balance sheet, in the same quarter that average claim severity is tracking toward record highs.

What Verisk found

The Q2 2026 numbers, as reported from Verisk's quarterly property report (Artemis.bm):

- **Claim assignment volumes fell to one of the lowest Q2 levels in recent years** — down 12.21% from Q2 2025 and 13.05% below the five-year average, extending a four-year decline.
- **The decline is concentrated in non-catastrophe claims:** non-CAT assignments down 18.65% vs. the five-year average, CAT assignments down just 4.29%.
- **CAT claims are now 43% of Q2 volume, up from 34% five years ago** — even though the number of PCS® designated events has held steady.
- **Average severity currently sits at \$17,085**, down 10.77% from Q2 2025 but only 2.88% below the five-year average — “tracking toward the highest quarterly level in recent years.”

That \$17,085 is almost certainly not the final number: severity data matures over two to three months as larger, more complex claims complete and return to Verisk's XactAnalysis system with higher replacement-cost values, so it generally appears artificially low during the “immature” period (Verisk Q1 2026 Quarterly Property Report, verisk.com). The Q1 precedent: initially reported at \$16,079, it has since matured 13.1% to \$18,185, beating last quarter's projection (Artemis.bm).

Apply a historical maturation rate and Q2 2026 lands near \$18,794; apply Q1's stronger rate and it lands above \$19,400 — the highest quarterly average in recent years. Both are projections and will move as data matures (Artemis.bm). The Q3 2025 edition showed the same pattern — assignments down 28.5% YoY while mature severity ran \$17,258–\$18,431 (Insurance Journal); the volume decline has run since 2023.

Why severity rises while designated events are stable

“Stable designated cat events” means the count of catastrophes formally designated by PCS — Verisk's Property Claim Services unit, whose serial-number designations are the industry's standard event definition — has held roughly steady. But event counts aren't cost exposure. Four mechanisms explain the divergence:

- 1. Mix shift.** Non-CAT volume collapsed (-18.65% vs. the five-year average) while CAT volume barely moved (-4.29%), making the remaining book mechanically more cat-heavy: 43% of Q2 2026 volume versus 34% five years ago (Artemis.bm). Catastrophe claims cost more, so the average rises with no exceptional single event.
- 2. Unit-cost inflation, led by labor.** Combined U.S. labor and material costs are up 4% year over year — and accelerated in each month of Q2 (Artemis.bm). Verisk's Q1 2026 report details it: total reconstruction costs +3.4% YoY, concrete masons +15.39% annually (led all trades), lumber -2.32%, fuel +42.70% (verisk.com).
- 3. Severity maturation.** The first number always understates reality: Q4 2025 climbed 10.15% post-report; Q1 2026 climbed 13.1% (verisk.com). Pricing off the immature number is pricing off a discount that doesn't exist.
- 4. Roof replacement severity, independent of storm counts.** Verisk's 2026 U.S. Roof Report: average residential roof replacement cost hit \$17,631 in 2025 — 33% above the prior four-year average — even though overall claims volume fell roughly 20%. Everyday wind and hail events, many below catastrophe thresholds, kept driving severity (GlobeNewswire via crweworld.com).

Two often-cited drivers — litigation and social inflation, and building-code changes — were not quantified in the sources behind this piece, so they are noted here rather than argued.

A related but different Verisk signal: the \$171B modeled-loss benchmark

Three weeks earlier, Verisk’s annual 2026 Global Modelled Catastrophe Losses Report set the industry’s average-year preparedness benchmark at **\$171 billion, up \$19 billion** — the highest estimate Verisk has reported (GlobeNewswire).

Do not conflate the two products:

- **The \$171B is a forward-looking modeled average annual loss (AAL)** — a cat-model output driven by exposure growth, rising reconstruction costs, and development in catastrophe-prone areas; a long-term benchmark, not a single-year forecast (Reinsurance News). It rose *despite* no U.S. hurricane landfalls for the first time in a decade.
- **The quarterly property report is backward-looking actual claims data** — real assignment volumes and real average severity from XactAnalysis. It shows the *mechanism* on the ground: fewer, more expensive, more cat-heavy claims.

Together: the modeled benchmark is rising because the unit economics of claims are rising — and the quarterly data proves it isn’t just model theory. As Verisk Catastrophe and Risk Solutions president Rob Newbold put it: “A quiet hurricane season can lead markets to respond as if risk has eased: rates soften, insurers keep more risk on their own books, and more capital competes to write new business. But 2025 reminds us that the underlying risk landscape has changed” (GlobeNewswire). For the sixth straight year, global insured catastrophe losses exceeded \$100 billion — driven not by hurricanes and earthquakes but by record wildfires and severe thunderstorm activity (Reinsurance News).

What this means for 2027 renewal pricing conversations

- **The “benign year” discount is a mirage.** Headline losses look tame, tempting buyers to expect softer pricing — but AAL just hit a record \$171B and actual severity is tracking toward record highs. Reinsurers anchoring on AAL won’t concede the risk got cheaper.
- **Watch attachments and aggregate deductibles, not just rate.** Severity plus inflation means industry losses could “ramp up more quickly than before” when major events occur — and ILS managers should weigh these trends “when negotiating key contract terms and features such as attachments and aggregate deductibles” (Artemis.bm). The 2027 negotiation is about structure as much as price.
- **Reserving is happening right now.** “Lower claim counts, a CAT heavier mix, and increasing severities are colliding with rising unit costs, especially labor, right as they finalize second half reserving, reinsurance, and vendor strategies” (Artemis.bm). Reserves set today flow into reinsurance demand and program structure at 1/1.

Regional note: the Midwest was the Q2 hotspot — Illinois and Ohio surged on severe convective storms and hail; Kansas and Iowa jumped 73% and 95% (Artemis.bm). Frequency perils keep generating widespread multi-community damage instead of single headline events — the pattern behind six straight \$100B+ years (Reinsurance News).

Sources: all figures cited inline, linked where they appear. Projections are labeled as projections; Verisk’s gated quarterly report figures are as reported by Artemis.bm.

Isaias and the Named-Storm Deductible Split

Source: <https://riskcoveragehub.com/isaias-and-the-named-storm-deductible-split/> Published: October 8, 2026

Hurricane Isaias became a coverage event at 10:00 a.m. CDT on October 8, 2026, not at landfall. National Hurricane Center Advisory Number 8 put maximum sustained winds at 85 mph and minimum central pressure at 975 millibars. The center was near 23.7 north, 90.2 west, about 380 miles south of the mouth of the Mississippi River, moving east-northeast at 10 mph. Landfall inside the warning area is forecast late Friday or early Saturday. A Hurricane Warning runs from Ocean Springs, Mississippi, to the Bay/Gulf County line in Florida. A Storm Surge Warning runs from the mouth of the Mississippi River to the Suwannee River. That designation, and those lines, are what policies will be read against.

Isaias is the first Atlantic hurricane of the 2026 season. Insurance Business reported on October 8 that months of calm had already pushed some policyholders toward larger deductibles and thinner flood buys, and that the gap between what people think a homeowners policy covers and what the form actually pays is the claim that will hurt. The advisory gives that argument a clock. Tropical-storm-force winds are expected in the warning area on Friday. Outside preparations are supposed to be finished today.

A named storm is a trigger, not a mood

Most Gulf property forms that use a hurricane or named-storm deductible define the event by a declaration from the National Weather Service or the National Hurricane Center, not by the adjuster's sense of the wind. Isaias is a hurricane on Advisory 8. If the policy keys the percentage deductible to a storm designated by name or number by NHC or NWS, this advisory is the document that starts the reading. If the form requires the storm to be a hurricane at landfall, the Friday night advisory is the one that matters, not this morning's 85 mph. Brokers who have not pulled the deductible endorsement before Friday are going to explain it from a screenshot.

The dollar difference is not theoretical. A flat all-other-perils deductible of \$2,500 or \$5,000 and a 2 percent or 5 percent named-storm deductible on a \$400,000 dwelling are different claims. Five percent of \$400,000 is \$20,000 before the first shingle is priced. On a commercial schedule the same percentage applied to total insured value, or applied separately per location, moves the retention by a factor of ten. Fannie Mae's selling guide, in the August 5, 2026 update to B7-3-02, caps the deductible for required perils on one-to-four-unit loans it buys at 5 percent of the coverage amount, including a separate windstorm or wildfire deductible. That is a loan-file ceiling, not a statement of what a surplus-lines or admitted Gulf policy charges. It is a reminder that 5 percent is already the outer band many lenders will accept, and that a storm week is when owners discover they are sitting on it.

Surge is not a wind claim

Advisory 8's surge figures are the exclusion. If peak surge meets high tide, water above ground could reach 5 to 7 feet from Ocean Springs to Indian Pass, including Mobile Bay; 4 to 6 feet from Indian Pass to the Steinhatchee River; 3 to 5 feet from the Mississippi mouth to Ocean Springs and from the Steinhatchee River to the Suwannee; and 2 to 4 feet from the Suwannee to Yankeetown, where a Storm Surge Watch is up. The National Hurricane Center describes that water as inundation moving inland from the shoreline. On a standard homeowners form, rising water is not wind. Flood, storm surge, and surface water sit in the water exclusion unless a flood policy or a specific endorsement puts them back.

The claim split on the same building will be boring and expensive. A lifted roof, a failed window, and wind-driven rain in the attic can be property causes of loss under a named-storm deductible. Water that came

over the slab from Mobile Bay or from a storm drain that reversed is a flood question. Anti-concurrent causation wording, where the form has it, will be cited on any loss where wind and surge arrived in the same night. The evidence that survives is the water line, the elevation certificate, and photographs taken before teardown. Adjusters will not reconstruct a 4-foot line from memory in November.

NFIP and private flood deductibles do not care that the wind deductible was already met. A building can owe a percentage named-storm deductible on the roof and a separate flood deductible on the ground floor. Contents limits on the flood policy are often lower than the homeowners contents limit the owner remembers. That mismatch is the quiet-season decision Insurance Business flagged on October 8: people cut premium in a year with no hurricane, then meet the first hurricane of the year with the cut still on the declarations page.

The rain band is a second policy reading

Wind and surge do not exhaust Advisory 8. From Friday through the weekend the center expects 3 to 7 inches from far southeastern Louisiana to the Florida Big Bend, with localized totals up to 15 inches in banding over the Panhandle and Big Bend. One to 3 inches, locally 5, can reach the Tennessee Valley and the Carolinas. River flooding is expected, some of it significant. Flash and urban flooding is possible, with considerable impacts flagged for southern Alabama, the Florida Panhandle, and southwestern Georgia.

That rain footprint is where backup-of-sewer, overflow, and surface-water wording gets read by people who are not in the Hurricane Warning. A sublimit of \$10,000 or \$25,000 for drain backup is a different product from a flood policy, and it often excludes water that enters from the surface rather than from a clogged line. Fifteen inches in a band does both. Reservation-of-rights letters on inland commercial accounts will turn on the point of entry: roof drain, dock, parking-lot sheet flow, or a river that left its bank on Saturday. Business-income forms that require a covered cause of loss at the described premises will not pay for a civil authority or off-premises power loss unless those endorsements are on the policy and the hour thresholds are met.

Time element starts before the check

Additional living expense and business interruption are limit-and-period problems, not deductibles, but Isaias sets the period up in a specific way. Hurricane-force winds extend only about 15 miles from the center. Tropical-storm-force winds extend about 90 miles. Power and road closures will outrun the eyewall. Isolated tornadoes are possible Friday night on the Alabama and Florida coasts and far southwestern Georgia, then Saturday across eastern Alabama, Georgia, and northern Florida. A 72-hour waiting period on civil authority or service-interruption coverage can consume the worst of the outage before the coverage responds. Owners who bought the buy-back and owners who did not will learn which stack they have when the generator log hits hour 40.

Mitigation does not wait on that reading. Once a building is safe to enter, pumping standing water and documenting the water line protects whatever coverage eventually applies. It does not create coverage the form excludes. Carriers will still ask for prompt notice. The practical sequence in the warning area is notice to the agent or carrier desk, a photo set, and extraction of water that is safe to remove, in that order if life safety allows. Debris removal sublimits and ordinance-or-law caps will matter more on older coastal stock than the headline deductible, especially where local officials require elevation or flood-vent repairs after the water leaves.

What to pull before Friday night

Four documents decide most of the first conversation. The declarations page, for the named-storm or hurricane deductible percentage and whether it applies per occurrence or per location. The water exclusion and any flood endorsement or separate flood policy, including the building and contents limits. The backup and discharge sublimit. The time-element endorsements, including waiting periods. None of those change because the 4:00 p.m. CDT advisory moves the cone ten miles. The cone does change which locations sit in the Hurricane Warning versus the Tropical Storm Warning flanks, from the Jefferson/Plaquemines Parish line to west of Ocean Springs, and from east of the Bay/Gulf County line to the Aucilla River.

Do not quote a modeled industry loss. Advisory 8 is a pre-landfall bulletin. There is no NHC damage estimate in it, and there should not be. The checkable figures are the ones on the bulletin: 85 mph, 975 millibars, a 15-mile hurricane-force wind field, a 90-mile tropical-storm-force field, surge of 5 to 7 feet in the eastern core segment, and rainfall of 3 to 7 inches with local amounts up to 15. Those are enough to sort wind from rising water and to tell a client why a quiet season did not retire the deductible they accepted in the spring.

From Cat Model Output to Underwriting Decisions

Source: <https://riskcoveragehub.com/using-cat-model-output-underwriting-pricing-decisions/> Published: October 1, 2026

Updated October 6, 2026.

Direct Answer: Catastrophe model output is not a price tag—it is an input to technical premium, retention, and limit choices that underwriters then reconcile with capacity, appetite, and market competition. When a carrier says “the model says \$40M,” your job is to trace that number through peril scope, return period, basis (ground-up versus gross), loadings, and reinsurance so you know what rate, deductible, and attachment actually follow—and when a documented override is in play.

You finished renewal prep with modeled average annual loss near \$2.1M and a 1-in-100 aggregate near \$40M. The indication lands 18% above last year’s rate on line and pushes a higher wind deductible. Both sides cite “the model.” This article walks the last mile: modeled loss to rate, retentions, technical versus charged premium, overrides, and questions to ask before you sign.

What the carrier’s number usually is—and what it is not

Before you argue rate, align on definitions. Most carrier conversations mix several distinct outputs:

- **Modeled average annual loss (AAL)**—expected loss per year before reinsurance, often ground-up or subject to a stated gross/net convention.
- **Occurrence or aggregate metrics**—for example a 1-in-100 occurrence loss or a 1-in-250 aggregate, each answering a different capital question.
- **Technical premium**—AAL plus loadings for expenses, cost of capital, and volatility, sometimes scaled to a target return period.
- **Charged premium**—what appears on the quote after competition, portfolio strategy, and underwriting judgment.

A broker slide that shows \$40M is rarely “your premium.” It is more often a tail metric the carrier uses for capacity, reinsurance purchase, or regulatory reporting—not an invoice line item. For how models are built, see the catastrophe modeling complete guide; this piece stays on the decision layer after those outputs exist.

From modeled loss to insurance rate: a worked path

Underwriters start from expected loss, then apply loadings. Illustrative arithmetic keeps the logic visible when carriers will not share every factor.

Step 1: Ground-up AAL to subject AAL

Take a \$500M total-insured-value schedule across ten coastal properties. The cat platform returns ground-up hurricane AAL of \$1.8M. Your policy has a \$100,000 per-occurrence deductible and a \$25M per-occurrence limit. The model’s “subject” AAL—the portion that would hit the policy after deductibles and limits—might fall to roughly \$1.4M after those terms are applied in the engine. Always confirm whether the carrier quoted ground-up or subject; mixing them is a common renewal dispute.

Step 2: AAL to technical rate on line

Technical premium often follows a form like:

Technical premium (Subject AAL ÷ limit) × load factor

Assume subject AAL \$1.4M, limit \$25M, load factor 1.65 (expense, cost of capital, and catastrophe volatility margin combined). Subject loss ratio on limit: $\$1.4\text{M} \div \$25\text{M} = 5.6\%$. Technical rate on line $5.6\% \times 1.65 = 9.2\%$. On a \$25M limit, technical premium \$2.3M before reinsurance credit.

Reinsurance changes the story: excess-of-loss cover can cut net expected loss while the gross quote still reflects gross model loss in tight markets. See reinsurance treaty structures explained for how layers interact with modeled aggregates.

Step 3: Technical versus charged premium

Charged premium adjusts technical premium for competition, relationship, and portfolio mix. A carrier may run 9.2% technical and quote 11.5% charged when cat capacity is scarce—or 7.8% to retain a strategic account. Reconstruct technical premium from shared assumptions rather than debating one percentage in isolation.

Retention and limit: translating tail output into structure

Model output drives not only rate but how much risk you keep. Underwriters watch tail metrics (100-year, 250-year) alongside AAL when setting deductibles, self-insured retentions, and policy limits.

Deductibles from modeled frequency and severity

Suppose wind deductibles are expressed as 2% of TIV per location, minimum \$250,000. Modeled results show that lowering the minimum to \$100,000 adds roughly \$180,000 to subject AAL on your schedule because more small and medium events pierce the retention. The underwriter proposes raising the minimum to \$500,000, which the model shows reduces subject AAL by about \$220,000 but leaves you retaining more ground-up loss on moderate events. The trade is arithmetic plus appetite: cash liquidity versus premium savings.

Limits and attachment from tail losses

If the 1-in-100 occurrence loss ground-up is \$85M and you buy \$50M primary with a \$50M excess layer attaching at \$50M, the model's gross loss at that return period informs whether the tower is credible or cosmetic. Portfolio managers also watch accumulation across locations; when multiple properties correlate in one event, location-level limits understate true exposure. For accumulation, PML, and reinsurance alignment, catastrophe portfolio management covers the upstream view—here, the underwriting question is whether your requested limit matches the modeled loss at the return period the carrier actually capitalizes to.

Technical premium versus market price at renewal

Market price is what clears in the placement window. Technical premium is what the carrier's internal models say is required for a target return. In soft markets, charged rates sit below technical for many accounts; in constrained cat markets, charged rates exceed technical for peak zones because supply caps bite before math does.

Document three numbers in your renewal file when possible: modeled subject AAL, implied technical rate on line, and quoted charged rate. The spread between technical and charged is information. Persistent charged-above-technical may signal sector stress or account-specific concerns (concentration, loss history, data quality flags). Charged-below-technical may signal competition—or an underwriter who expects remedial terms later if loss experience deteriorates.

Modeling costs affect how much independent checking you can afford; see catastrophe modeling costs and fees. A lighter run on the same exposure snapshot helps test whether the carrier's subject AAL is in the same ballpark.

When underwriters override the model—and how they should document it

Models are tools, not statutes. Overrides are normal when empirical data or engineering judgment materially disagrees with catalog assumptions—provided they are scoped, authorized, and recorded.

Common override triggers

- **Site-specific mitigation**—hurricane strapping, flood barriers, or wildfire defensible space not fully captured in geocoding or building attributes.
- **Book versus file mismatch**—submission TIV or construction class differs from modeled values after inspection.
- **Peril scope differences**—flood excluded on policy but still appearing in a blended “nat cat” dashboard export.
- **Recent loss experience**—credible local events that stress-test vulnerability beyond average years.
- **Portfolio strategy**—temporary appetite expansion or contraction not encoded in vendor defaults.

Overrides should state direction, magnitude, peril, dates, and approver. Weak documentation surfaces at audit and claim time. Ask for a summary when quoted terms move materially from a transparent technical build—especially deductible or sublimit changes tied to “model update.”

Questions to ask when the carrier cites “the model”

Polite, specific questions beat generic pushback. Use them in broker strategy calls and underwriter meetings:

- Which metric are you citing—AAL, occurrence loss at 100 years, aggregate at 250 years—and on what basis (ground-up, gross, net of reinsurance)?
- Which event set and version date were used, and did peril scope match our policy (wind-only versus full nat cat)?
- What load factor converts AAL to technical premium for our limit and layer?
- How much of the rate change is model movement versus loading, expense, or market adjustment?
- What retention change produces how much modeled premium delta on our schedule?
- Was any underwriting override applied? If yes, what attribute or experience drove it?

When one peril dominates your schedule, hurricane, earthquake, and wildfire peril analysis adds context; you still only need consistent definitions to negotiate rate and retention from the carrier’s authorized output.

How buyers should use model output at renewal

Treat your model export and the carrier’s model as two witnesses, not one verdict. Align inputs first: same locations, values, construction, deductibles, and limits. Then compare subject AAL and one agreed tail metric. If gaps persist, split “data disagreement” from “terms disagreement.” Data gaps get fixed with surveys and engineering; terms gaps get structured with deductibles, limits, and alternative markets.

Build a one-page renewal logic map: subject AAL, implied technical rate, quoted rate, retention deltas, and open override questions. Share it early so marketing targets carriers whose appetite matches your modeled profile.

When catalogs refresh, ask what changed on your schedule—frequency, severity, or indexing—not only the headline percent change. See climate risk pricing and catastrophe model updates for background; at renewal, tie term movements to pieces you can verify.

Present ranges from your own vendor run and ask the carrier to reconcile to their authorized platform. How major vendors differ is covered in RMS, AIR, and Verisk in catastrophe modeling; here the objective is a decision-ready bridge to bound terms.

Closing the gap: dashboard to deal

The model says \$40M at a stated return period; your job is to translate that into whether \$2.3M technical premium for a \$25M limit feels fair at a 9.2% rate on line, whether a \$500,000 wind minimum is worth \$220,000 of modeled premium relief, and whether overrides are documented when the quote diverges from transparent math. Carriers that explain that chain earn faster binds; buyers who understand it avoid signing structures priced on metrics they never agreed to. That is the last mile between the dashboard and the deal—and it is where catastrophe modeling pays off in dollars, not slides.

FAQ

What is the difference between technical premium and charged premium?

Technical premium is the amount an insurer’s pricing framework says is needed to cover modeled expected loss plus planned loadings for expenses, cost of capital, and catastrophe volatility. Charged premium is what the insurer actually quotes or binds after competitive pressure, account strategy, and underwriting judgment. The

two often differ; understanding both lets you see whether a renewal move is driven by model loss change or by market and discretion.

How does modeled loss become an insurance rate?

Underwriters usually convert subject average annual loss into a loss ratio by dividing AAL by the policy limit, then multiply by a load factor that reflects expenses and return requirements. That yields a technical rate on line; charged rate adjusts from there. Always confirm whether AAL is ground-up or subject to your deductibles and limits before comparing rates year over year.

How are retentions set from model output?

Carriers simulate how deductibles, self-insured retentions, and co-insurance change the loss that attaches to the policy. Each retention option produces a different subject AAL and tail loss, which maps to a different premium. Underwriters balance modeled premium savings against your liquidity and the frequency of events that will pierce higher retentions.

When do underwriters override the model?

Overrides are common when verified mitigation, corrected exposure data, peril scope mismatches, or credible loss experience are not fully reflected in the vendor run. Effective overrides are directional, quantified where possible, approved at the right authority level, and documented with peril and date context. Buyers should ask for a plain-language summary when terms diverge materially from a transparent technical build.

What should I ask when the carrier cites the model?

Ask which metric and return period they cite, ground-up versus subject basis, event set and version date, load factors used for technical premium, and how much of the change is model versus market adjustment. Ask whether overrides were applied and how a proposed deductible or limit change shifts modeled premium on your schedule. Request consistency with the policy's peril scope and your submitted values.

How should a buyer use model output at renewal?

Reconcile inputs with the carrier, compare subject AAL and at least one tail metric on agreed terms, and separate data fixes from structure negotiations. Build a short logic map from modeled loss to implied technical rate to quoted premium and retention choices. Use independent or broker runs to sanity-check subject AAL, and press for documentation when charged premium diverges from explainable technical premium.

Model Blending and Sensitivity Testing in Cat Risk

Source: <https://riskcoveragehub.com/model-blending-sensitivity-testing-cat-risk/> Published: October 1, 2026

Updated October 6, 2026.

Direct Answer: Model blending is the disciplined practice of combining outputs from two or more catastrophe models—or the same model under explicit alternative assumptions—into a single view of loss that no single vendor run can justify on its own. Professionals weight models by peril and portfolio fit, stress vulnerability curves, event rates, and financial terms, and document every choice so underwriters, reinsurers, and auditors can see why the renewal number is what it is—not a black-box average.

Two public facts, not a new blending method. Florida’s 2023 law lets property insurers average two or more commission-approved hurricane models; before that, a filing had to pick one accepted model (Insurance Journal, June 12, 2025). Separately, Aon’s 2025 survey of its own clients and prospects — a survey, not a census, and not a license-dollar market share — found that about half of respondents license no vendor and lean on a broker. Moody’s RMS and Verisk are the two predominant licenses. Verisk is wider among U.S. respondents. Moody’s is preferred in UK/EMEA. Reinsurance placements in that survey are built on one of those two or a blend of them, because that is what the reinsurer will re-run. No percentages are printed here: the document does not state how many people answered.

When your catastrophe report shows one expected loss, one probable maximum loss, or one tail metric, the practical question is rarely which model is right. It is whether any one model, frozen at one version and one set of assumptions, should carry the full weight of a placement, a retention decision, or a board narrative. Model blending and sensitivity testing make combination and challenge explicit: how blends differ from naive averaging, how weights vary by peril, what to stress, and how to record the work so the dashboard number has a defensible story. For how cat modeling fits the wider renewal stack, see the catastrophe modeling complete guide for 2026.

Why One Modeled Number Is Never Enough

Portfolios sit at the intersection of exposure quality, peril mix, and contract language. No single module excels on every peril and region. Blending does not fix bad exposure; it acknowledges that each model is a structured opinion on physics, statistics, and damage—not a tape measure.

Risk managers and CFOs often receive one modeled AAL or 1-in-250 year loss on a slide. Blending reconciles parallel model worlds before the market does opaquely; sensitivity testing shows which assumptions move pricing, attachment, or capital. That pairing is the discipline behind the dashboard.

What Model Blending Means in Practice

Model blending assigns weights to two or more modeled loss distributions—or key metrics—and produces a composite view across vendors, versions, or approved overlays. Outputs include blended AAL, merged exceedance curves, or fixed return-period PML. Weights must be reproducible and tied to peril rationale.

Reinsurance panels often expect proof you did not blindly adopt the lowest or highest vendor quote. This workflow is not EP-curve reading or vendor shootouts; it is the arithmetic and governance of combination itself.

Metrics Versus Full Curves

Metric blending forms a weighted sum of AAL or 250-year PML from two runs. **Curve blending** merges exceedance probabilities at each loss level, then reads off metrics. Metric blending is simpler to audit; curve blending preserves tail shape when weights are stable across return periods. If weights change by return period—more weight to Model A below retention and Model B above—you are curve blending, not averaging two PMLs at one point.

Blending Versus Simple Averaging

Simple averaging—50/50 between two AALs, or the mean of three vendor PMLs—looks neutral and rarely is. Equal weights imply equal credibility across perils, regions, and model generations, which almost never

matches reality. Averaging also ignores correlation: two models built on similar catalogs and vulnerability libraries do not provide independent information; their mean can overstate certainty.

Weighted blending starts from explicit hypotheses—higher weight where catalog treatment, extratropical behavior, or occupancy mapping fits your book. A judgment layer is allowed only if defined, not smuggled in as rounding.

Worked Example: Weighted AAL Versus a Straight Mean

Take a \$500M total-insured-value schedule with two vendor AALs for the same peril and territory: Vendor 1 reports \$4.2M; Vendor 2 reports \$6.8M. A straight average is $(\$4.2M + \$6.8M) / 2 = \$5.5M$. Suppose your peril committee assigns 65% weight to Vendor 1 and 35% to Vendor 2 after a side-by-side on exposure mapping and flood sublimit handling. The blended AAL is $(0.65 \times \$4.2M) + (0.35 \times \$6.8M) = \$2.73M + \$2.38M = \$5.11M$. The \$390K gap versus the simple mean is documented relative trust, not noise. If sensitivity testing shows Vendor 2's AAL moves more than \$1M when wind deductibles follow manuscript wording exactly, you may shift weights at renewal.

Weighting Models by Peril

Peril-specific weights are the norm in multi-peril programs. Hurricane, earthquake, severe convective storm, wildfire, and flood draw on different modules and validation paths. A weight set for Florida wind is inappropriate for California earthquake without re-argument. Many organizations maintain a **weight matrix**: rows are perils or peril-region buckets, columns are models or scenarios, cells sum to 100% within each row.

Weights connect to historical fit, independence of catalog and vulnerability logic, update stability, and reinsurer pricing alignment. Document carve-outs when licensing limits leave part of the book single-sourced.

Regional Splits Within a Peril

Large portfolios sub-weight within hurricane: higher weight where storm surge and BI waiting periods match coastal wording, lower weight inland where models agree within a few percent. Earthquake blends split strike-slip versus subduction zones. Wildfire blends may down-weight a vendor where wildland-urban interface coding was mass-updated mid-term. Collapsing that heterogeneity into one national weight without narrative is what auditors push back on. Align weights with where models structurally disagree using hurricane, earthquake, and wildfire catastrophe peril analysis as peril context, not as a substitute for your matrix.

Sensitivity Testing: What to Stress

Sensitivity testing varies defensible inputs and records how key metrics move. Blending sets the central view; sensitivity defines error bars and escalation triggers. Standard domains are vulnerability, event rates, and financial terms—the levers underwriters manipulate when comparing options.

Vulnerability, Event Rates, and Financial Terms

Stress vulnerability curves, catalog frequency, and financial terms together. Upper-bound vulnerability for a dominant occupancy can move PML more than a small weight shift between vendors. Rate tests use approved scalars or discrete activity views to test stability, not to pick a climate storyline. Terms tests rerun the same events with manuscript-faithful deductibles, coinsurance, BI waiting periods, and sublimits—where gross AAL agrees but net PML diverges, terms—not physics—often drive the gap.

Worked Example: Sensitivity on a Single PML

Suppose blended 250-year hurricane PML is \$48M on a \$500M TIV program with a \$5M per-occurrence deductible. Vulnerability high case: +18% on wind-driven damage for key occupancies yields PML about \$56.6M. Event rate +15% on the catalog yields PML about \$55.2M. A diagnostic run that applies deductibles per building instead of per occurrence might show \$41M—terms sensitivity larger than vendor spread. For an excess layer attaching at \$50M, vulnerability and rate move the metric across the attachment; deductible mechanics determine whether the layer responds in the model.

Governance, Documentation, and External Review

Blending without workflow becomes spreadsheet folklore. A minimal path: lock exposure snapshot and model versions; run baselines; agree peril-level weights with recorded sign-off from risk, underwriting, and analytics; compute blended metrics; execute sensitivity tests; publish a blend memo tied to renewal. Budget sometimes limits vendor runs—document single-sourced perils honestly; see catastrophe modeling costs and who pays for fee allocation, not weight logic.

Underwriters need weights by peril, central metrics, and sensitivity on attachment or named storm sublimits; flag any gap versus bound deductibles. Tie metrics to treaty form—aggregates versus cat excess—using reinsurance treaty structures: quota share, excess of loss, and aggregate. Memos should list snapshot date, versions, weight table, merge rule, blended AAL and tail metrics, sensitivity deltas, approvers, and internal versus submission use.

Auditors and reinsurers want repeatable weight rationale, synchronized terms, and change logs. Sensitivity should bracket the base case. Align narrative with catastrophe portfolio management, accumulation, and PML for reinsurance, and note whether vendor updates refreshed the blend or weights stayed fixed to isolate exposure growth—see climate risk pricing and catastrophe model updates in 2026.

Pitfalls and Platform Context

Do not treat correlated models as independent, blend gross with net, or let weights drift after sign-off. Side-by-side diagnostics inform weights without becoming a vendor contest; platform background sits in how RMS, AIR, and Verisk approaches quantify catastrophe risk, while your memo cites weights and outcomes only.

Before Renewal Sign-Off

Confirm peril-level credibility, sensitivity on vulnerability, event rates, and financial terms, and reconstructability of the central metric from the memo. Then the dashboard figure reflects governed blending—not a single oracle run.

FAQ

What does model blending mean in catastrophe risk?

Model blending is the explicit combination of loss results from two or more catastrophe model runs into one coherent view, using defined weights or merge rules rather than treating a single run as definitive. The inputs can be different vendors, different versions, or the same engine under approved alternative assumptions. Outputs may include blended average annual loss, exceedance curves, or return-period metrics such as PML.

The purpose is to reflect peril-specific credibility and portfolio fit while keeping the method transparent for renewal and audit.

How is blending different from simple averaging?

Simple averaging assigns equal importance to every input, which assumes equal credibility across perils, regions, and model designs that rarely match reality. Weighted blending applies peril-specific or region-specific weights tied to documented criteria such as historical fit, exposure handling, or treaty relevance. Averaging also treats highly correlated models as if they were independent, which can narrow apparent uncertainty falsely. Blending makes those judgments visible instead of hiding them inside a neutral-looking mean.

How should models be weighted by peril?

Build a matrix where each peril or peril-region row has weights summing to 100% across models or scenarios used for that row. Tie weights to explicit factors: quality of exposure mapping for that peril, stability across recent updates, independence of catalog and vulnerability assumptions, and alignment with reinsurer pricing for the same peril. Avoid one national weight if coastal wind, inland storm, or seismic zones behave differently in your book. Document carve-outs where only one model ran because of data or licensing limits.

What does sensitivity testing cover?

Sensitivity testing varies defensible inputs and records how key metrics move. Standard domains are vulnerability or damage functions, event rates or catalog frequency treatments, and financial terms including deductibles, coinsurance, BI waiting periods, and sublimits. Teams usually report percent changes to AAL and to tail metrics such as 100-year or 250-year loss against a documented base blend. The goal is to show which assumptions can change placement decisions, not to replace the central blend with the worst case.

How should a model blend be documented?

Publish a blend memo with exposure snapshot date, model names and versions, peril-level weight tables, merge formulas, blended central metrics, and a short sensitivity summary with named approvers. State which metrics are for internal planning versus carrier or reinsurer submission and confirm financial terms match bound or quoted manuscripts. When weights change from prior renewals, log the reason. Store files with version control so mid-term edits do not overwrite the record shared at renewal.

What do auditors and reinsurers expect to see in a blend?

They expect repeatability: weights justified by criteria, synchronized policy terms, and a clear trail when assumptions change year over year. They look for sensitivity results that bracket the central estimate rather than a single unchallenged base case. Reinsurers compare your metrics to their own models; documentation explains intentional differences without appearing arbitrary. They also expect honesty about single-model perils, known data gaps, and whether major vendor updates were incorporated or deliberately held fixed to isolate exposure growth.

Climate-Conditioned Catastrophe Model Catalogs

Source: <https://riskcoveragehub.com/climate-conditioned-catastrophe-model-catalogs/> Published: October 1, 2026

Updated October 6, 2026.

Direct Answer: A climate-conditioned catastrophe catalog replaces—or sits beside—the practice of sampling future losses from events anchored mainly in historical storms, quakes, and fires. Vendors shift frequencies, intensities, and footprints using defined warming scenarios so modeled hurricane, flood, and wildfire tails can rise even when recent loss experience looks quiet. At renewal, treat the output as a labeled scenario layer on your baseline run, not a single authoritative premium—and state which catalog, pathway, and horizon the report assumes.

Your renewal appendix already depends on an event catalog whether anyone used that term or not. The engine draws from a library of plausible catastrophes, applies geocoded values, and rolls up losses. For decades that library inherited the climate embedded in history—observed events plus stochastic resampling for gaps. Climate conditioning breaks that default: the library is rebuilt or reweighted so draws reflect a stated future baseline, not only the last fifty years. That shift moves return-period labels, reinsurance sizing, and the gap between modeled loss and quoted terms. The overview at catastrophe modeling complete guide for 2026 places catalogs inside the wider stack; here we stay on what changes when warming scenarios enter the event set and how to read it without turning renewal into politics.

What the catalog controls in your report

The catalog is the supply of catastrophes the model can replay—tracks, ruptures, ignition patterns, and modifiers such as surge or pluvial add-ons. It decides which hurricane strengths can reach your Gulf sites and how often fire weather hits your wildland-adjacent assets. Historical catalogs extend the observational record with synthetic completion; stochastic methods perturb tracks so you are not limited to replaying named storms exactly. Even then, category mix and drought years inherit past climate. Climate-conditioned catalogs adjust or regenerate that hazard layer for a forward time slice and scenario, often leaving exposure and vulnerability modules unchanged between runs.

Why boards hear that the past is not the catalog

Modeled medians and lived tails diverged in years when coastal wind, flood, and urban-edge wildfire losses stacked up. Models always lag exposure growth and inflation, but conditioning targets a different lag: assuming the next thirty years draw from the same hazard mix as the last thirty. Physical risk teams call that non-stationarity; insurers frame it as capital and pricing. You need not settle every projection dispute to see why lenders ask for forward-looking evidence—and why a historical-only run may no longer be the only story in the room.

Historical sets versus climate-conditioned sets

A historical stochastic catalog answers: given past hazard statistics, what loss distribution emerges for this portfolio? A climate-conditioned catalog asks: given a defined warming scenario and horizon, how should event frequencies and severities change before the same loss math runs? You may receive two tables for identical locations. Historical might show a one-percent aggregate exceedance probability (AEP) loss of \$48 million; climate-conditioned might show \$61 million when Gulf and Atlantic landfall weights shift toward higher categories or when wildfire draws gain weight in corridors history undercounted. Neither table is the model alone—they are catalog choices through the same engine.

Capture three metadata fields on every climate run: scenario label (often tied to a representative concentration or shared socioeconomic pathway), horizon year (2035 versus 2050), and whether the vendor replaced the baseline catalog or supplied a parallel view. Brokers sometimes circulate the higher number without those fields; underwriting alignment depends on you pushing them back into the thread.

How vendors inject warming into event libraries

Pipelines differ by vendor, but most combine regional climate signals with catastrophe physics so events stay plausible.

Reweighting and regeneration

One pattern keeps historical track templates but changes draw weights. Illustration: a hurricane catalog has one thousand landfalling Category 2-class events and two hundred Category 3-class events. Conditioning might shift weights to eight hundred and three hundred twenty while holding shapes stable—expected loss rises because stronger bins draw more often, not because someone multiplied every loss by 1.2. Another pattern regenerates tracks under warmer oceans and shifted steering, moving landfall latitude and rain fields. That matters for compound wind-flood outcomes where a small track change moves large insured values.

Horizons, ensembles, and scope

Deliverables often expose near-term and mid-century slices rather than one blended future. Some workflows average multiple climate models; others pair a central estimate with a high-sensitivity envelope. Ask whether “climate-adjusted” means full catalog replacement, a licensed module, or an optional overlay—scope drives catastrophe modeling fees and who pays. Peril rollout is uneven; hurricane and flood conditioning lead, while earthquake views may emphasize tectonic time dependence rather than greenhouse pathways.

Tail estimates and worked numbers

Tails surface in retention and reinsurance. Take \$500 million total insured value in coastal industrial wind, historical catalog: 100-year loss \$75 million, 250-year loss \$120 million. Return periods are catalog-defined annual exceedance probabilities (1 percent and 0.4 percent), not literal calendar waits. A 2050 central climate catalog on the same exposure might show \$92 million at 100-year (+23 percent) and \$158 million at 250-year (+32 percent). Upper bins often gain more relative weight than median bins—the “shape steepens.”

Average annual loss (AAL) can move less. Historical AAL \$4.0 million might become \$4.6 million (+15 percent) while the 250-year point rises +32 percent. Toy arithmetic: if 90 percent of AAL comes from events below \$10 million and 10 percent from the tail, raising tail contributions 50 percent while holding the lower band flat gives AAL $\$4.0 \text{ million} \times (0.9 + 0.1 \times 1.5) = \$4.6 \text{ million}$. Real portfolios add peril correlation and terms; the pattern still explains why portfolio accumulation, PML, and reinsurance focus on EP tails after conditioning. Run at least two scenarios where appetite is set—central versus lower pathway might span \$140 million to \$158 million on the same 250-year metric.

Regulatory push and pushback

Supervisors and disclosure regimes press insurers to show how climate may shift portfolios, which feeds forward-looking analytics in filings and enterprise risk reports. Corporate buyers feel it through capacity, diligence questionnaires, and rate filings even when no single statute names climate catalogs. Pushback comes

from actuarial caution about opaque adjustments and from jurisdictions resisting mandated scenarios in retail ratemaking, citing uncertainty and political conflict. Keep records factual: scenario, horizon, module scope, and whether conditioned views drove pricing or only reporting. Peril mechanics sit in hurricane, earthquake, and wildfire peril analysis without merging physics into catalog policy here.

Renewal conversations that stay on terms

Lead with year-over-year deltas: historical AAL and 250-year loss, then conditioned counterparts if produced. Ask which catalog generated the submission number. Tie language to retention, sublimits, and tower attachment: “Historical 100-year wind is \$75 million; carrier 2050 central view is \$92 million—confirm the reinsurance program sized on historical, conditioned, or the higher of the two.” Mixed catalogs make markets look cheaper or harsher for the wrong reason; align peril scope with the carrier and note when flood weights change but other perils do not. Market pricing context beyond catalogs appears in climate risk pricing and catastrophe model updates. File a one-page catalog memo—vendor, scenarios spelled out, geographies most affected, percent change at 100-year and 250-year—so CFOs can explain tail rises in quiet claim years.

Before you bind

Fix geocode and values first. Request paired historical and conditioned runs with identical financial terms when budget allows. Stress retention against the higher tail, not AAL alone. When reinsurance renews, state whether treaties referenced historical or conditioned metrics; attachment disputes have followed mismatches. Treaty forms interact with modeled tails in reinsurance treaty structures. Major platforms have piloted conditioned modules by peril and region—ask which event set the carrier used. See how major modeling platforms quantify risk for consistent questions across sources.

Conditioned catalogs do not replace mitigation; they change which synthetic futures your balance sheet rehearses. Dual-view reporting—history for continuity, scenarios for stress—is the disciplined path.

FAQ

What is a climate-conditioned catalog?

A climate-conditioned catalog is a catastrophe model event library whose storm, flood, fire, or other peril frequencies and severities have been adjusted to reflect a defined future climate state and time horizon. It is built so the modeling engine can draw events that are consistent with warming pathways—not only with the historical observational record. Your loss metrics still flow from the same exposure and vulnerability data; the hazard supply changes. Treat it as a scenario-based event set with explicit metadata rather than a silent default.

How does a climate-conditioned catalog differ from a historical event set?

A historical event set derives its hazard statistics from past observations and established extensions, then uses stochastic methods to fill gaps. A climate-conditioned set changes those statistics—or regenerates events—to match a forward climate scenario while keeping physical plausibility constraints. Side-by-side runs on the same portfolio often show similar median outcomes with wider gaps at high return periods. The difference is the question being asked: past hazard mix versus plausible future hazard mix under stated assumptions.

What does climate conditioning do to modeled tail loss?

It often increases high return period losses more than it increases average annual loss because warming adjustments frequently add weight to severe event bins. A portfolio might show a mid-teens percent AAL lift with a larger percent lift at the 100-year or 250-year point, depending on peril and region. Tail movement is scenario-dependent; lower and central pathways diverge. Use tails for retention and reinsurance sizing only together with the scenario label that produced them.

Which vendors offer climate-conditioned views?

Leading commercial catastrophe modeling vendors have released or are piloting climate-conditioned modules or event sets for selected perils and territories, with rollout tied to licensing and regional filings. Offerings change by model version and geography; no single public list stays current without your vendor's release notes. Ask your broker or consultant which modules your carrier actually used in pricing. Require written confirmation of the event set name and scenario parameters rather than relying on marketing summaries.

What is the regulatory angle on climate-conditioned models?

Regulators and disclosure frameworks increasingly expect insurers to demonstrate how climate change may affect future losses, which pushes adoption of forward-looking analytics in filings and enterprise risk reports. At the same time, some jurisdictions resist mandating climate scenarios in retail ratemaking, citing uncertainty and political disagreement. Corporate buyers feel the effects through carrier capacity, disclosure requests, and lender diligence rather than one uniform rule naming climate catalogs. Keep filings factual: scenario, horizon, and whether conditioned views affected pricing or only reporting.

How should I discuss climate-conditioned output with my broker at renewal?

Bring a short comparison table: historical versus conditioned AAL and key return period losses, with scenario and horizon spelled out. Ask which catalog the submitting carrier used and whether reinsurance followed the same choice. Tie numbers to decisions—retention, sublimits, and tower attachment—instead of debating climate policy. Request that markets quote on consistent assumptions so spreads reflect terms, not hidden catalog swaps. File the metadata with your renewal packet so next year's comparison is honest.

Secondary Perils in Cat Modeling: SCS, Flood, and Wildfire

Source: <https://riskcoveragehub.com/secondary-perils-cat-modeling-scs-flood-wildfire/> Published: October 1, 2026

Updated October 6, 2026.

Direct Answer: Severe convective storm, flood, and wildfire often dominate modeled and actual property loss outside hurricane and earthquake zones. Models treat them as high-frequency, dispersed perils with attritional tails—not single peak events—so annual averages, protection gaps, and deductibles must be read against the peril that actually drives loss.

Your renewal pack still leads with hurricane PML and earthquake scenario loss, but the lines that move are hail, convective wind, pluvial and fluvial flood, and wildfire—including smoke and ember that never match your wildfire sublimit wording. For multi-state portfolios, that is a coverage design problem. Below: why primary-peril framing undercounted these drivers, how models treat them, and how to test deductibles and sublimits against protection-gap arithmetic.

Model releases, 2024–2025 (not a ranking)

- May 28, 2025: Moody’s released RiskLink / RiskBrowser Version 25, including North Atlantic Hurricane Models Version 25. Florida accepted Version 25 on June 12, 2025, valid until November 1, 2027.
- June 12, 2025 and June 19, 2025: Florida accepted Verisk Hurricane Model for the United States Version 3.0.0 on Touchstone 2024A and Touchstone 2025.
- July 24, 2025: California Department of Insurance: first completed wildfire-model review is the Verisk Wildfire Model. CDI was still reviewing Karen Clark & Company and Moody’s RMS.
- October 2, 2024: Moody’s North America Wildfire HD Version 2.0.
- June 2025: Verisk Touchstone / Touchstone Re 2025 (13.0): U.S. severe thunderstorm update, and inland flood for the UK and Republic of Ireland.
- December 10, 2025: Moody’s North America severe convective storm HD models. Loss figures in that launch are Moody’s claims, not an industry audit.

Why “Secondary” Stopped Meaning Small

A secondary peril is any cat peril that is not the account’s designated primary driver—usually hurricane or earthquake. SCS, inland flood, and wildfire stay labeled secondary on schedules even when they produce the largest single-year loss. Legacy benchmarks optimized for low-frequency, high-severity peaks; SCS and flood produce many moderate events nationwide, and wildfire concentrates in the wildland–urban interface. Over a decade those shapes often dominate expected loss versus a primary tail that never realized.

Reports still center wind tiers and quake zones while hail swaths sit in appendices. If secondary-peril contribution is material but the tower prices a 1-in-250 hurricane, layers are misaligned. The catastrophe modeling complete guide for 2026 covers the full stack; this piece stays on SCS, flood, and wildfire loss and program structure.

Severe Convective Storm: Frequency Over Peak

SCS bundles hail, tornado, straight-line wind, and related precipitation. Models simulate large catalogs of smaller footprints, not one mega-storm. SCS average annual loss scales with roof area, age, material, and spread—not a single landfall parameter.

Worked example: hail as an attritional layer

Take a \$500M TIV schedule in central and southern states with aged built-up roofs. Modeled SCS AAL 0.12% of TIV (\$600,000) before deductibles:

$$0.0012 \times \$500,000,000 = \$600,000 \text{ modeled SCS AAL.}$$

With a 2% per-occurrence deductible, only loss above that layer nets to you; gross SCS AAL still drives roof inventory and ACV questions. One regional outbreak can hit hundreds of locations—attritional in claims, fully represented in deep SCS catalogs.

Why did SCS losses explode in industry discourse? Exposure in hail-prone metros, larger locations, and higher values per square foot raise ground-up loss without any storm-climate change. Expanded SCS catalogs and finer hail resolution mean current runs attribute more loss than older versions on the same footprint. When

comparing runs, ask whether the delta is catalog, vulnerability, or TIV—not only weather narrative.

Flood: Multiple Hydrologies, One Policy Gap

Flood is several hydrologies: fluvial, pluvial, surge, and special scenarios in different modules. Secondary-peril flood here means inland rain and river loss outside a named hurricane—often excluded or sublimited on property forms.

Protection-gap arithmetic

Ground-up flood versus insured flood after terms defines the protection gap. On a \$200M campus, modeled 100-year ground-up flood might be \$18M while property terms assume \$0 insured flood without a private limit or NFIP.

Economic flood can read \$18M while insured PML is zero under exclusions. Add a \$5M flood sublimit: insured loss = min(ground-up, sublimit) after deductible, so \$18M ground-up caps at \$5M insured and \$13M protection gap. Reproduce that min() by location when quoting facultative or parametric cover.

Flood is path-dependent on rainfall, soil moisture, drainage, and pad elevation—not zip code alone. Uncertainty is wide at address level; neighbors can diverge by feet of elevation. Blanket sublimits without geospatial review misstate insured loss. For wildfire alongside hurricane and earthquake context, see hurricane, earthquake, and wildfire catastrophe peril analysis; flood still demands hydrologic resolution.

Wildfire: Flame, Ember, and Smoke as Separate Mechanisms

Wildfire modules separate flame, ember ignition, and smoke where architecture allows. Flame hugs the perimeter; ember extends downwind; smoke can drive BI far from the burn. Policy wildfire sublimits may not match how claims code smoke versus fire—align model peril buckets to endorsements before comparing tails.

Concentration versus spread

Take \$300M TIV across 50 WUI locations: wildfire AAL 0.08% (\$240,000) but a 250-year ground-up cluster scenario of \$45M. A \$10M sublimit caps insured tail at \$10M; \$35M stays gap. Earthquake deductibles rarely trigger on ember; wind deductibles may not apply when loss codes as fire.

Models ingest fuel, slope, access, and mitigation credits where available. Ember spotting beyond the flame front drives divergent loss on opposite sides of a ridge. Treat year-over-year wildfire delta as exposure in the WUI plus catalog change, not automatic retention movement.

How Secondary-Peril Modeling Differs From Primary-Peril Modeling

Primary engines stress sparse tail events; secondary modules stress catalog breadth and seasonality. SCS curves look smoother below 100-year returns; insured flood can sit at zero until sublimits bind; wildfire tails stair-step by cluster. Use AAL for SCS negotiations; pair flood and wildfire AAL with tail scenarios at sublimit kneepoints.

Data prep differs: SCS needs roof and geocode precision; flood needs elevation and hydrology; wildfire needs fuel proximity. A few hundred feet of geocode error can flip flood or ember loss. Secondary perils punish address hygiene more than coarse quake zones.

Platform choice affects peril splits—see catastrophe modeling across RMS, AIR, and Verisk for context—but read your own export first. Fee allocation for extra layers is in catastrophe modeling costs and who pays; misaligned sublimits usually cost more than the run.

From Model Output to Deductibles and Sublimits

Secondary perils stack peril-specific triggers. On \$400M TIV, suppose 2% all-risk deductible, \$15M flood sublimit, \$10M wildfire sublimit west, and full limits for SCS.

Illustrative hail: \$22M ground-up SCS; 2% on \$180M affected TIV = \$3.6M deductible; \$18.4M net if limits apply. Pluvial flood \$12M ground-up walks another path through flood sublimit and water deductibles. Each secondary peril uses different triggers in the same tower.

Hail is not windstorm on every form; flood sublimits without geospatial matching cap the wrong sites; wildfire sublimits may exclude smoke while BI pays. Treaties may treat SCS as attritional unless a complex threshold is met—AAL can sit below cat reinstatement while aggregates erode retention. See reinsurance treaty structures including quota share and excess-of-loss aggregate for how aggregates absorb many SCS years versus one wildfire tail.

Portfolio View: Accumulation Without a Landfall

Hail swaths, river basins, and fire complexes create accumulation without a landfall cone. Portfolio reviews that only cluster hurricanes miss SCS within hail swaths and flood by watershed. Tie secondary-peril aggregate metrics to reinsurance, not only named-storm PML—see catastrophe portfolio management for accumulation and PML.

Model-update cycles refresh secondary-peril catalogs often; see climate risk pricing and catastrophe model updates in 2026 for how that feeds pricing dialogue. Use updates to re-test sublimits, not to auto-raise retention.

What to Do Before Renewal

Request peril-disaggregated SCS, flood, and wildfire output with policy terms in the engine: ground-up and insured AAL plus one tail metric each. QA geocode and elevation on top locations per peril. Map deductibles and sublimits to model peril codes. Recompute protection gap as ground-up minus insured at the same return period. Secondary in name only, these perils are often primary in frequency and expected loss—align limits to the peril that sends the check.

FAQ

What counts as a secondary peril in catastrophe modeling?

In catastrophe modeling, a secondary peril is any modeled peril that is not the portfolio's designated primary peril for rating and reporting—most often hurricane or earthquake. Severe convective storm, inland flood, and wildfire are routinely labeled secondary even when they drive the largest insured or economic losses in a given year. The label describes convention and report layout, not low importance. Models still simulate secondary perils with full catalogs, vulnerability functions, and financial modules. Your job is to read their metrics on the same footing as primary peaks when they affect limits and retentions.

Why did severe convective storm losses explode?

SCS losses grew because exposure increased in hail and convective-wind regions, values per square foot rose, and roof inventories aged into higher vulnerability bands. Model catalogs and hail-size resolution also improved, so current runs attribute more loss to SCS than older platform versions for the same footprint. Large regional outbreaks produce many simultaneous location claims that aggregate above retentions. That pattern is distinct from a single hurricane landfall but can exceed annual cat budgets. Separating exposure growth from catalog change keeps renewal debates honest.

How does secondary-peril modeling differ from primary-peril modeling?

Primary-peril modeling emphasizes rare, high-severity events with deep tail extrapolation for wind and quake. Secondary-peril modeling emphasizes many smaller to mid-size events, seasonal catalogs, and fine spatial footprints for hail, rain, and fire. Average annual loss is often more informative for SCS; flood and wildfire need tail review with sublimits applied. Secondary perils depend heavily on location attributes—roof, elevation, fuel—rather than coarse zones alone. Insured loss outputs require explicit policy terms because economic and insured curves diverge more than for all-risk wind on standard forms.

What makes flood so hard to model?

Flood loss hinges on local hydrology, elevation relative to drainage paths, rainfall intensity, and antecedent conditions—not on a single county flag. Pluvial flooding can occur far from mapped flood zones; fluvial flooding depends on upstream behavior. Small geocode or elevation errors can change loss dramatically. Policy exclusions and sublimits often zero out insured flood while economic loss remains, producing a wide protection gap. Models propagate uncertainty bands that brokers should not collapse to a single point without site review.

How do wildfire models treat ember versus flame damage?

Wildfire models typically simulate direct flame exposure at the fire front separately from ember ignition downwind and, where supported, smoke or ash impacts. Flame damage correlates with fuel continuity and structure spacing at the perimeter; ember damage extends into subdivisions that never see direct fire. Vulnerability curves differ by ignition pathway—ember losses often start at vulnerabilities such as vents, roofs, and decks. Policy language may not mirror these buckets, so align model peril codes with endorsements before comparing insured tail loss to sublimits.

What do secondary perils mean for my deductibles and sublimits?

Secondary perils mean deductibles and sublimits must be tested peril by peril, not only against hurricane or earthquake scenarios. SCS may erode retention through many mid-size hits; flood and wildfire sublimits cap tails below modeled ground-up loss and create protection gaps. Deductible wording for wind, water, and fire may not match how events are coded in claims or in the model. Reproduce $\min(\text{limit}, \text{ground-up loss})$ after deductibles for each secondary peril at AAL and at a relevant tail. That arithmetic shows whether your tower matches the peril driving modeled loss.

Moody's RMS and Verisk: Two Catastrophe-Model Lineages, Not Three

Source: <https://riskcoveragehub.com/catastrophe-model-vendors-rms-air-verisk-differences/> Published: October 1, 2026

Updated October 6, 2026.

Direct answer: There are not three catastrophe-model vendors named RMS, AIR, and Verisk. There are two dominant lineages. Moody's RMS is Risk Management Solutions, bought from DMGT; that deal closed September 15, 2021. AIR is the other lineage: Applied Insurance Research, folded under the Verisk name in January 2022. Trade press dated November 14, 2025 says Verisk renamed Extreme Event Solutions to Catastrophe & Risk Solutions. "Verisk's own models" means that ex-AIR stack. ISO rating data, PCS loss estimates, and Xactware are sister products, not a third stochastic engine.

A renewal still moves when the carrier's license, the version, or a second view changes the tail. What changed is the map of who sells the model. For the four-module workflow, use the catastrophe modeling complete guide. This page is only the naming map and where to read the rest.

The two lineages

Moody's RMS. Risk Management Solutions was founded in 1988. Moody's agreed on August 5, 2021 to buy it from Daily Mail and General Trust and closed the purchase on September 15, 2021. Models and software now sit inside Moody's: RiskLink, Risk Modeler on the Intelligent Risk Platform, and a newer high-definition (HD) framework beside older detailed loss models. Not every RMS model is HD. Florida still certifies a versioned North Atlantic hurricane model on RiskLink and Risk Modeler. Moody's said on May 28, 2025 that no RiskLink / RiskBrowser model updates were then scheduled for 2026, with a caveat that urgent science could land first on the cloud platform.

AIR, now Verisk. Karen Clark founded Applied Insurance Research in 1987. Insurance Services Office bought AIR in May 2002. Verisk was not the buyer that year; Verisk Analytics later became the holding company around the ISO businesses. On January 31, 2022 Artemis reported that Verisk was rebranding AIR Worldwide under the parent for its extreme-event business. On November 14, 2025 Captive International reported the division name moving from Extreme Event Solutions to Catastrophe & Risk Solutions, with the experts and workflows unchanged. Platforms in market: Touchstone and Touchstone Re. Verisk's June 2025 release notes say Touchstone / Touchstone Re 2025 (13.0) is the last full release of those products, with Verisk Synergy Studio planned for June 2026 to replace both. As of October 2, 2026 this page does not treat Synergy Studio as launched. Support URLs in the 13.0 notes still pointed at air-worldwide.com. That is one science stack with a new label, not a vendor competing with itself.

Deal-announcement lines ("400 models," "underlie the industry," Artemis's early-2022 cat-bond modeller share) are seller or tracker claims for those dates. They are not 2026 license share. That 2022 cat-bond figure should not be treated as current share.

What is not a third model

ISO statistical and rating data, PCS industry-loss estimates, and Xactware claims estimation sit in the same corporate family as the ex-AIR catastrophe models. They do not add a third hazard engine next to Moody's RMS and Verisk. Verisk also markets Model Exchange as access to third-party catastrophe models. That is a shelf of other people's models, not "Verisk peril analytics" as a separate stochastic family. A submission that cites "Verisk" and "AIR" on the same account is usually naming one lineage twice.

Distinct other families

These are not the big two, and they are not a false three-way either.

- **Karen Clark & Company.** Clark left AIR. KCC sells its own reference models on RiskInsight. Florida accepted KCC's US Hurricane Reference Model 5.0 on June 11, 2025. "Most accurate" language on the vendor site is marketing.
- **Cotality, formerly CoreLogic.** CoreLogic announced the global rebrand on March 24, 2025. That release is a corporate identity change, not a claim that the hurricane model was new science that day. Florida's accepted-model list still labels the organization "CoreLogic/Cotality" and accepted its Florida Hurricane Model 2025 on June 11, 2025.
- **Aon Impact Forecasting,** on the ELEMENTS platform. Named by the Florida commission and in Aon's own 2025 client survey.
- **Florida Public Hurricane Loss Model** (Florida International University). Version 8.3 was accepted June 13, 2025.

Florida's list of models acceptable under the 2023 hurricane standards, revised March 3, 2026, also includes Moody's North Atlantic Hurricane Models 25.0 and Verisk's U.S. Hurricane Model 3.0.0 (Touchstone 2024A and Touchstone 2025). Those hurricane acceptances on that list expire November 1, 2027. Insurance Journal on June 12, 2025 still named Applied Research Associates among models the commission had accepted. The March 3, 2026 list under the 2023 standards does not include ARA. Whether an older certificate is still in force was not re-checked. The acceptance list is not a license census.

Why the two numbers still disagree

Both lineages use the same four-stage skeleton: hazard, vulnerability, exposure, financial. What differs is how each stage is computed and how large and time-aware the event set is. Moody's describes HD models as a temporal simulation on a uniform grid, with damage at coverage level, and says HD event sets can run up to eight million stochastic events. That eight-million figure is a vendor ceiling in a 2023 explainer, not a count for every peril. Verisk's June 2025 notes describe updates to hazard, vulnerability, and the stochastic event set, plus a changed financial model in Touchstone 13.0.

A Version 25 curve and a Touchstone hurricane 3.0.0 curve can share axis labels (AAL, occurrence EP, aggregate EP, PML) and still be different distributions. Read the definitions on how to read an EP curve. This is not a scorecard. No source reviewed for this revision publishes a neutral, same-portfolio delta, so none is invented here, and this page does not repeat unverified lines about which region leans on which vendor.

Public revenue market share was not found. The only usable adoption evidence behind this revision is Aon's 2025 survey of its own clients and prospects, plus regulatory acceptance lists. Read it as a survey, not a census. About half of respondents license no vendor and lean on a broker. Moody's RMS and Verisk are the two predominant licenses. Verisk is wider among U.S. respondents. Moody's is preferred in UK/EMEA. No percentages are printed here on purpose: the survey does not state a respondent count, its footer marks the PDF confidential even though the URL is public, and a one-point gap sits between the prose and a bar chart.

For how a desk actually blends those views, use model blending and sensitivity testing. For why severe convective storm, flood, and wildfire now move the program, use secondary perils. Dated releases belong on that page, not as a second explainer here.

What the desk should record

Record vendor, product (RiskLink, Risk Modeler, HD, Touchstone, Touchstone Re, RiskInsight, ELEMENTS, or the Florida public model), version, and run date. A model-year change can move technical premium and the layer a reinsurer will sell even when total insured value did not change. Florida is the clearest public version of that rule: residential rate filings use a commission-accepted model, and a 2023 Florida law lets property insurers average two or more of those accepted models.

Platform lock-in is part of model risk in 2026. Moody's is still supporting several RiskLink generations at once. Verisk has told clients Touchstone 13.0 is the last full release and that Synergy Studio is the planned replacement. Which buttons the underwriter presses is not a footnote.

The practical choice for most carriers is which of the two dominant licenses, which version, and whether a second view is blended in: the other major vendor, KCC, Cotality, Impact Forecasting, or an internal overlay. There is still no universal "right" model. The operative one is the one that binds capacity.

FAQ

Are RMS, AIR, and Verisk three vendors?

No. Moody's RMS is one lineage. AIR is Verisk's catastrophe-model lineage, rebranded under the parent in January 2022 and, in trade press from November 2025, housed in Catastrophe & Risk Solutions. ISO, PCS, and Xactware are not a third stochastic model.

Who else sells a distinct model?

Karen Clark & Company, Cotality (CoreLogic until the March 24, 2025 rebrand), Aon's Impact Forecasting, and the Florida Public Hurricane Loss Model are separate families. Florida's March 3, 2026 acceptance list under the 2023 hurricane standards names those alongside Moody's and Verisk. Acceptance is not the same as license share.

Why do two accepted models disagree on the same book?

Catalogs, vulnerability curves, and financial engines differ. The labels (AAL, EP, PML) do not. Compare the same metric, return period, peril scope, and financial perspective, then read the EP-curve page before arguing the tail.

What should a renewal file down?

Vendor, platform, version, and run date; whether the number is a license run, a broker-supplied view, or a blend; and which accepted Florida model (or average of accepted models) a residential filing actually used.

Sources for the dates above: Moody's (August 5 and September 15, 2021; May 28 and June 13, 2025), Insurance Journal (May 15, 2002 and June 12, 2025), Artemis (January 31, 2022), Captive International (November 14, 2025), Verisk's June 2025 Touchstone 13.0 release summary, Cotality (March 24, 2025), the Florida commission list revised March 3, 2026, and Aon's 2025 Global Catastrophe Risk Management Survey (public URL; confidential legend; a survey, not a census).

How to Read an EP Curve: Occurrence vs Aggregate Loss

Source: <https://riskcoveragehub.com/how-to-read-ep-curve-occurrence-aggregate-cat-modeling/> Published: October 1, 2026

Updated October 6, 2026.

Direct Answer: An exceedance probability curve plots how likely it is that losses will reach or exceed a dollar amount in a defined period—usually one year—not a single “expected” number. Occurrence EP answers the worst single event; aggregate EP answers total loss from all events combined. The 100-year point is a loss level with roughly a 1% annual chance of being equaled or exceeded, not a forecast that you will lose exactly that amount every century. Reading the curve correctly is what turns a modeling report from a black box into a renewal, reinsurance, and capital conversation you can defend.

Your catastrophe modeling report probably arrived as tables, maps, and at least one chart with a gently sloping line and unfamiliar axis labels. Underwriters and brokers may cite a “250-year loss” or “1% EP” as if it were one fixed fact. The curve is the fact: it encodes modeled outcomes, including the rare tail that drives reinsurance and retentions. This article stays on how to read that curve—occurrence versus aggregate, exceedance probability on the axes, return-period points, and why the far right matters. For the wider workflow, see the catastrophe modeling complete guide for 2026.

What your report is showing on the axes

Most vendor and broker EP exhibits use loss on the horizontal axis (sometimes logarithmic) and exceedance probability on the vertical axis, or they swap axes and keep the same relationship. Exceedance probability is the chance that losses in the modeled year will meet or exceed a given threshold. Reports may instead show return period, the approximate inverse of that annual probability. A 1% exceedance probability corresponds to a 100-year return period when the curve is on a per-annual basis.

Read the footnotes before you read the line. The curve is meaningless without peril scope, financial perspective (ground-up, gross, or net of reinsurance), and exposure basis (values, limits, deductibles, business interruption). A ground-up occurrence EP curve for Florida wind is not comparable to a net aggregate EP curve for all perils nationally unless you normalize those choices. Peril context appears in our hurricane, earthquake, and wildfire catastrophe peril analysis.

Probability versus return period

If the vertical axis is labeled “EP” or “exceedance probability,” a point at 2% means there is a 2% modeled chance per year that losses reach at least the loss on the horizontal axis at that height. If the axis is labeled “return period in years,” the same point might read 50 years, because $1 \div 0.02 = 50$. Brokers often speak in return periods; model output often speaks in EP. They are two labels on the same idea.

Do not treat return period as a countdown timer. A 100-year loss is not “due” after ninety-nine quiet years. It is statistical shorthand for a low-probability, high-severity zone on the curve. Multiple years can exceed that threshold in a row, or you can go decades without touching it, and both outcomes remain consistent with the EP framing.

Occurrence EP versus aggregate EP

Occurrence exceedance probability describes the distribution of the largest single event loss in the year (confirm how the report defines the period). Aggregate exceedance probability describes the distribution of the sum of all event losses in the year. For a schedule with one dominant peak per season, occurrence and aggregate can look similar at high return periods. For a wide footprint with many events, aggregate EP at moderate

thresholds often sits to the right of occurrence EP because several smaller events can add up to a total no single occurrence reaches.

Your report says “occurrence EP” when the question is how bad one hurricane, one quake, or one wildfire cluster could be. It says “aggregate EP” when the question is whether annual retained loss could exhaust an aggregate retention, erode an aggregate stop, or trigger an aggregate cat bond index. Treaty language follows the same split: per-occurrence limits respond to occurrence loss; annual aggregates respond to aggregate loss. Our overview of reinsurance treaty structures including quota share, excess of loss, and aggregate maps those mechanics to the curves you see.

A simple numeric contrast

Take illustrative annual event outcomes for one location year, not model output: Event A costs \$40M, Event B costs \$35M, and no other events apply. Occurrence loss is \$40M (the maximum single event). Aggregate loss is \$75M (\$40M + \$35M). On an occurrence EP curve, the point at \$40M reflects the probability of any single event reaching \$40M. On an aggregate EP curve, the point at \$75M reflects the probability the sum of events reaches \$75M. Per-occurrence retention tracks the occurrence curve at attachment; an annual aggregate deductible tracks the aggregate curve at that sum.

What the 100-year point really means

When someone says “the model’s 100-year loss is \$X,” they mean the loss level on the EP curve where annual exceedance probability is about 1%, or return period is about 100 years, under the report’s peril, perspective, and catalog assumptions. It is a quantile of the modeled distribution, not a budget line item you should expect every year. Average annual loss (AAL) sits far to the left on the same curve—typical years—not in the tail.

Reports rarely print exactly 1.00% EP; you may see 0.95% or 1.05% at the quoted dollar amount. Small changes in uncertainty treatment or loss amplification assumptions can move the 100-year dollar materially. Vendor methodology differs, as summarized in how RMS, AIR, and Verisk approaches quantify catastrophe risk; your job is to read the curve and footnotes in hand.

Why the tail wags the tower

EP curves are flat on the left and steep on the right. Moving from the 60-year to the 30-year point often changes the dollar amount modestly. Moving from the 250-year toward the 100-year region can change it dramatically because you are sampling the sparse tail of the simulated catalog. Reinsurance attachments and aggregate triggers often live in that tail. A modest shift in modeled 250-year occurrence loss can mean a layer clears your program or leaves a gap.

The body of the curve sets AAL and budgeted cat load; the tail sets capacity and solvency stress. Portfolio roll-ups compound tails through correlation and concentration; see catastrophe portfolio management, accumulation, PML, and reinsurance for schedule-level context. On a single EP chart, the tail still dominates whenever someone asks about “worst case” or “1-in-100.”

Occurrence tail versus aggregate tail

For multi-state wind and hail schedules, aggregate tail loss can exceed occurrence tail loss at the same return period because the model sums events across the footprint. For earthquake with one dominant rupture scenario, occurrence and aggregate tails may nearly coincide. When your report shows only one curve, confirm which

one it is before you attach it to a treaty clause.

Worked reading: \$500M total insured value

Assume a simplified commercial property schedule with \$500M total insured value across coastal wind-exposed locations, ground-up perspective, hurricane peril only, occurrence EP, one-year horizon. The report's curve (illustrative, not a vendor quote) might read approximately as follows:

At 10-year (10% EP), modeled occurrence loss might sit near \$25M—enough to matter for deductibles and lower tower layers. At 50-year (2% EP), perhaps \$95M. At 100-year (1% EP), perhaps \$140M. At 250-year (0.4% EP), perhaps \$220M. These numbers are invented for pedagogy; your report will differ. The exercise is how to use them.

Suppose your physical damage deductible is \$5M per occurrence and you carry \$100M excess of \$5M per occurrence. At 50-year occurrence loss of \$95M, retained loss is \$5M and the layer pays \$90M. At 100-year loss of \$140M, the layer pays \$100M and \$35M remains uninsured or flows to a higher layer if you bought one. The EP curve gives a ladder of thresholds so you can ask whether each rung matches a contract.

Switch to aggregate EP on the same schedule. If the 100-year aggregate point is \$180M while the 100-year occurrence point is \$140M, a treaty with a \$150M annual aggregate retention could be penetrated by aggregate loss even when no single occurrence exceeds \$140M—multiple events in one season. Phrase it for your file: 100-year occurrence \$140M, 100-year aggregate \$180M means single-event capital peaks lower than annual summed stress.

Check whether the curve is before or after reinsurance. A gross occurrence EP curve at \$140M at 100-year does not tell you net retention until you apply the program. Net EP curves shift left when reinsurance is efficient and can kink where layers attach. If the report only shows ground-up, reconstruct net with your layer chart or ask the broker for a net view.

The “your report says X” decoder

“AAL is \$8M.” Expected value of annual loss across simulated years—center of mass of the EP curve, not the tail. Useful for pricing; weak for limit adequacy.

“1% occurrence EP loss is \$140M.” Same as 100-year occurrence in annual framing; one-event severity at that quantile.

“250-year aggregate is \$310M.” Sum-of-events view at 0.4% annual exceedance; use for aggregate covers and some cat bond triggers.

“EP curve shifted higher at 100-year after the update.” Tail repricing, not necessarily proportional change to AAL or every layer premium. Ask what changed: exposure, vulnerability, or event set. Broader update economics sit in climate risk pricing and catastrophe model updates in 2026; your curve still demands line-by-line reading.

“PML” on a broker summary. Often a return-period loss with internal definitions. Match the acronym to the EP point on the chart; do not assume PML equals a number from a different curve type. When tables and charts disagree, trust the documented export and reconcile before bind.

Using EP curves in renewal dialogue

Bring the curve, not only the headline 100-year figure. Mark deductibles, layer attachments, and co-participation on the loss axis for the relevant occurrence or aggregate EP. Ask where the net curve crosses those marks at 50-year, 100-year, and 250-year. That aligns risk management, broking, and finance on what is priced versus retained.

For buying decisions, compare marginal tail: incremental loss between 100-year and 250-year points relative to premium for the covering layer beats debating AAL alone. Modeling fees are a separate procurement topic—see catastrophe modeling costs, fees, and who pays.

Common chart variants and axis traps

Some exhibits plot EP on a logarithmic scale to compress the tail. “Non-exceedance” flips the vertical axis: that curve rises left to right while an exceedance curve falls. Reading the wrong orientation inverts tail intuition.

Multi-year EP curves exist for bond tenors. A five-year EP is not the same as a one-year EP at the same return-period label. Align the horizon with the contract before you cite a point.

FAQ

What does an EP curve actually show?

An EP curve shows the modeled probability that losses in a defined period will meet or exceed each loss amount on the chart, given the perils, exposure, and financial perspective in the report. It is a full distribution summary, not a single forecast. Each point pairs a dollar threshold with an exceedance probability or return period. You use it to see typical years on the left and rare tail outcomes on the right.

What is the difference between occurrence and aggregate exceedance probability?

Occurrence EP describes the largest single event loss in the period, while aggregate EP describes the sum of all event losses in the period. A year with two major events can produce a high aggregate loss with a lower occurrence loss. Treaty structures mirror the distinction: per-occurrence limits respond to occurrence EP, and annual aggregates respond to aggregate EP.

What does the 100-year loss point on an EP curve mean?

The 100-year point is the loss level at roughly 1% annual exceedance probability under the report’s definitions. It means the model assigns about a one-in-one-hundred chance per year that losses reach or exceed that amount, not that the loss repeats on a calendar schedule. It is a tail quantile for planning and reinsurance, not expected annual spend.

Why does the EP curve matter more than average annual loss?

AAL summarizes the average of many simulated years and sits on the low-probability side of the curve, which aligns with pricing load but not with limit adequacy or insolvency stress. Reinsurance attachments, retentions, and board tail questions live at 100-year, 250-year, and beyond. The curve shows those regions; AAL alone does not.

How do EP curves set reinsurance attachment points?

Teams overlay occurrence or aggregate EP points with proposed attachments and limits to see where net loss remains at selected return periods. If 100-year occurrence loss pierces a proposed retention, the tower or terms adjust until net EP at key points fit appetite. Aggregate EP supports annual aggregate deductibles and aggregate stop-loss structures.

What are the most common misreads of EP output?

People treat return period as a due date, confuse occurrence with aggregate, compare curves with different peril or net-versus-gross definitions, and cite AAL when the question is tail adequacy. Another frequent error is ignoring footnotes on catalog, epistemic uncertainty, and exposure date. Always match the curve type to the contract clause before you bind.

What Catastrophe Modeling Costs: Fees, Licenses, and Who Pays in 2026

Source: <https://riskcoveragehub.com/catastrophe-modeling-costs-fees-who-pays/> Published: October 1, 2026

Updated October 6, 2026.

Direct Answer

Catastrophe modeling spend in 2026 usually splits into enterprise software licenses, per-run or bureau fees, and optional data or analytics add-ons. Enterprise licenses for the major vendor platforms commonly run from the tens of thousands of dollars per year into seven figures annually, depending on peril modules, geographic coverage, user seats, and whether analytics or regulatory filing tools are included. Per-run pricing—typical when a carrier, MGA, or broker submits a portfolio to a modeling bureau rather than running models in-house—often lands in a per-location or per-policy band from a few dollars to several tens of dollars per run for standard property lines, with surcharges for detailed engineering data or secondary perils. Primary insurers and large reinsurers typically pay the largest license bills; brokers and MGAs more often pay bureau or subscription access fees and pass costs through to clients or embed them in program economics. For what that money buys, see the sections below; for how models produce loss estimates, see Catastrophe Modeling: The Complete Guide.

Licensing models: enterprise seats, per-run fees, and data add-ons

Vendor economics cluster around three patterns: full enterprise licenses, transactional per-run or bureau pricing, and recurring data-feed subscriptions layered on either.

Enterprise and seat licenses. AIR, RMS, and Verisk (among others) sell multi-year enterprise agreements to carriers and reinsurers that include core modeling engines, selected peril and territory modules, a defined number of concurrent or named users, and baseline maintenance that covers model version releases. Pricing is quoted privately; contracts rarely publish list prices. Buyers negotiate modules (for example, North Atlantic hurricane, U.S. earthquake, European windstorm), server or cloud deployment rights, and whether exposure management and portfolio analytics sit in the same entitlement.

Per-run and bureau pricing. When an organization lacks a full license—or needs independent verification—it sends exposures to a licensed bureau or to a counterparty with enterprise access. The fee is tied to volume: policies, locations, or total insured value processed. MGAs, regional carriers, and program administrators often use this model for peak underwriting seasons or delegated specialty programs. Some reinsurers and brokers

operate shared platforms and charge participants per submission or per account roll-up.

Data and analytics add-ons. High-resolution hazard layers, geocoding enrichment, building-level replacement-cost feeds, and climate-conditioned or forward-looking hazard scenarios typically carry separate annual fees. These extend the core license rather than replacing it. Teams tracking climate risk pricing and catastrophe model updates often budget add-ons alongside maintenance because regulatory and rating discussions increasingly reference supplemental views beyond default vendor catalogs.

Realistic cost ranges as of 2026

Public pricing is scarce; the ranges below reflect common industry reporting and procurement patterns, not any single vendor quote. Treat them as planning bands and confirm numbers in RFPs.

Enterprise licenses for a mid-size U.S. property carrier with multi-peril domestic coverage and a modest user count typically fall somewhere between roughly \$500,000 and \$2 million per year in total annual contract value when maintenance, core modules, and standard support are bundled—though smaller regional footprints with fewer perils can land below that band, and global multiline groups with many seats and custom analytics can exceed it by a wide margin. Entry-level or limited-peril entitlements for specialty insurers or reinsurance sidecars sometimes start in the low hundreds of thousands annually.

Per-run bureau economics for standard commercial property schedules often cluster from about \$3 to \$25 per location or policy equivalent for basic geocoding and primary-peril views, with higher tiers when engineering-style building attributes or flood, surge, or secondary-peril modules are invoked. Portfolio roll-ups for reinsurance pricing may be quoted as flat project fees from five figures into six figures depending on account count and iteration rounds.

Data-feed and scenario add-ons commonly add five figures to low six figures per year per feed or region bundle. Implementation, cloud surcharges, and excess API use can add 10–30 percent to first-year spend for shops leaving legacy servers.

None of these figures should be read as list prices from AIR, RMS, or Verisk; they are composite ‘what buyers report’ ranges for budgeting and board questions about the catastrophe modeling fee.

Who pays across the value chain

Primary carriers and mutuals. Balance-sheet insurers with material property books almost always carry the enterprise license and the largest internal modeling teams. The catastrophe modeling fee sits in IT and analytics budgets, sometimes split with underwriting when exposure management platforms are shared. Regulatory filing and rate-support work is a major internal consumer of that spend.

Reinsurers. Global reinsurers mirror carrier license structures and often maintain parallel vendor stacks for treaty and facultative pricing. They pay for deep portfolio analytics tied to catastrophe portfolio management, accumulation, and PML workflows. Cedents rarely reimburse those license costs directly; they show up in reinsurance quotes and capacity allocation described in the reinsurance professional guide.

Brokers and intermediaries. Large brokers may hold enterprise access for placement analytics, catastrophe bond structuring, or client reporting, or they may purchase per-account runs from bureaus and embed fees in brokerage or advisory charges. Wholesale and specialty brokers on delegated authority programs often pass

per-run costs to MGAs or program carriers contractually.

MGAs and insurtechs. These entities rarely justify a full enterprise license at launch. They rely on bureau per-run pricing, rented seats on a carrier's or reinsurer's platform, or white-label analytics from capacity providers. As premium volume grows, unit economics push some toward partial licenses or dedicated cloud tenant deals. Property program mechanics and limit structures that drive exposure counts are covered in the property insurance professional guide.

What the fee actually buys

A license or run fee is not a single 'model run' button. Buyers receive a bundle of intellectual property, software, and services.

Hazard modules supply event catalogs, footprints, and intensity fields for agreed perils and regions.

Vulnerability modules translate hazard at the asset level into damage ratios by construction and occupancy; the engineering assumptions are vendor-specific, and curve-level detail lives in the parent catastrophe modeling guide rather than here. **Financial modules** apply policy terms—limits, deductibles, coinsurance, and reinsurance placeholders—to produce insured loss estimates suitable for underwriting and reporting, aligned with concepts in limits, deductibles, and coinsurance practice.

Maintenance and model updates deliver new event sets, vulnerability revisions, and bug fixes on vendor cadences—often annually for major peril refreshes. Contract language defines whether updates are mandatory (and potentially repriced) or optional modules.

Support and training range from help-desk access to dedicated client teams for Tier 1 accounts. **Regulatory and rating filings** sometimes require vendor-certified output formats or documentation; enterprise deals may include filing assistance or pre-built regulatory views where vendors offer them.

Per-run bureau fees typically buy a defined peril set, standard financial treatment, and a report artifact—not unlimited reruns, custom vulnerability, or full portfolio optimization tooling unless separately scoped.

How buyers evaluate ROI

Finance and underwriting leaders rarely ask whether modeling is 'worth it' in the abstract; they tie spend to decisions with dollar consequences.

Rate adequacy and portfolio mix. Modeled loss costs anchor technical rates for property lines. If modeled average annual loss or tail metrics drift from experience, product teams adjust class plans, moratoriums, or geographic appetite before regulatory scrutiny intensifies.

Reinsurance purchasing. Modeled PML and TVaR-style metrics feed treaty attachment, retention, and pricing discussions. A modeling fee that improves visibility into zone overlap or sub-peril gaps can outweigh its cost in a single renewal if retention or collateral requirements move favorably.

Accumulation and growth controls. Exposure management against county, zip, or custom zones prevents silent aggregation in fast-growing books. ROI shows up as avoided surprise post-event and as better alignment with capital models.

Counterparty and investor credibility. Reinsurers, regulators, and rating agencies expect consistent vendor use or documented equivalents. Underinvestment surfaces in model risk management findings or slower approvals for new products.

Modeling without an enterprise license

Smaller carriers, MGAs, and startups have practical paths that stay within economics proportionate to premium.

Bureau per-run contracts with minimum commits spread cost across the underwriting year. **Capacity-provider platforms** let delegated underwriters model on the reinsurer's or carrier's entitlement, sometimes at no direct license fee but with exclusivity or margin expectations. **Consulting and analytics firms** offer project-based modeling for rate filings or one-off portfolio reviews. Parametric and index-based products reduce bespoke cat modeling for some risks but do not eliminate exposure analytics for the rest of the book.

Hybrid approaches—light internal analytics plus selective bureau runs for peak zones—are common until written premium and regulatory footprint justify a full RFP to AIR, RMS, Verisk, or alternative vendors. Model risk governance still applies: document assumptions, version IDs, and who performed the run, regardless of fee structure.

Frequently asked questions

How much does catastrophe modeling cost for a mid-size insurer in 2026?

Mid-size U.S. property carriers with multi-peril domestic programs often report total annual vendor spend—including enterprise license, maintenance, and common add-ons—in a band from roughly half a million to two million dollars, though footprint, seat count, and module mix move outcomes above or below that range. Per-run bureau users at smaller scale may spend far less in absolute dollars but more per policy at low volume.

What is the difference between an enterprise cat model license and per-run pricing?

An enterprise license grants ongoing access to vendor software, selected peril modules, updates, and internal run capacity subject to contract limits. Per-run pricing charges each time exposures are processed on someone else's license, typical for bureaus and outsourced analytics. Enterprise suits high repeat volume; per-run suits seasonal or low-volume books.

Who typically pays for AIR, RMS, or Verisk catastrophe models?

Primary insurers and large reinsurers pay the largest enterprise licenses. Brokers may pay for analytics platforms or pass through bureau run fees. MGAs and insurtechs usually pay per-run fees or use a capacity provider's platform rather than holding their own enterprise agreement.

Are catastrophe model prices published publicly?

No. Major vendors quote privately based on modules, regions, users, and term length. Industry articles and RFP benchmarks provide planning ranges only; procurement teams should treat any single number as illustrative until a vendor or bureau proposal is in hand.

What does a catastrophe modeling fee include besides software?

Contracts typically bundle hazard and vulnerability intellectual property, financial loss engines, model version updates under maintenance, technical support, and defined deployment rights. Optional feeds, climate or

scenario overlays, extra API volume, and professional services are usually priced separately.

Can a small MGA get cat modeling without a seven-figure license?

Yes. Bureaus, reinsurer-hosted platforms, and project-based consultants offer per-run or subscription access scaled to volume. Costs accrue per location or policy and rise with optional perils and data enrichment, but absolute spend stays below full enterprise tiers for most startup and mid-growth program administrators.

Atlantic ACE 92% Below Normal: Quiet CAT Does Not Soften Wind Deductibles

Source:

<https://riskcoveragehub.com/atlantic-ace-92-below-normal-quiet-cat-does-not-soften-wind-deductibles/>

Published: September 24, 2026

Updated October 6, 2026.

Direct Answer: Through October 1, 2026, the Atlantic basin had eight named tropical storms, zero hurricanes, and accumulated cyclone energy (ACE) of 10.8—about 89 percent below the 1991–2020 normal for the date, per the National Hurricane Center seasonal table. That quiet stretch does not commute percentage hurricane or named-storm wind deductibles on bound policies, and it is not a mid-term lever to cut catastrophe load.

Brokers are already hearing the sentence: “No hurricane by October — so why is the wind deductible still a percentage of Coverage A?” The answer is in the ACE number and the filed form, not the headline count of named storms.

NHC data through 21 UTC October 1, 2026, showed eight named Atlantic storms—Arthur through Hanna—with no system reaching the 64-knot hurricane threshold. Basin ACE stood at 10.8, 89 percent below the 1991–2020 normal for that date. Mid-September the same table sat near 6.3 ACE with six names; late-September additions Fay (peak 60 knots), Gonzalo, and Hanna pushed activity without producing a hurricane. Fay’s peak equates to about 69 mph, still below the 74-mph hurricane line. The season entered October without an Atlantic hurricane for the first time in the modern reconnaissance and satellite era—a longer wait than the prior late-first-hurricane marks in 2002 and 2013.

NOAA climatology still puts the average first Atlantic hurricane around August 11 and expects four hurricanes by September 16. NOAA’s 2026 August outlook boxed the year at 7–13 named storms and 2–6 hurricanes. What has verified so far is a weak, shear-limited pattern—not a reason for carriers to rewrite percentage wind deductibles mid-term.

ACE is the catastrophe unit. Named storms are not.

Treaty and cat-model conversation runs on ACE, average annual loss (AAL), and probable maximum loss (PML), not on how many times the NHC issued a name. A 40-knot two-day storm barely registers. The events that move coastal property results are long-tracked hurricanes that stack surge, inland flood, and demand surge across several states. 10.8 ACE through the historical peak means those events have not booked. It does not mean the peril left the schedule or the deductible paragraph.

Edouard still reached the northern Gulf coast as a tropical storm in early September—a live property file with a named-storm trigger, not grounds for a mid-term endorsement deleting the hurricane deductible. The NHC

seasonal summary shows zero direct deaths and zero U.S. damage in the aggregate table so far; working-layer water and wind files can still run under ordinary deductibles without rewriting CAT terms.

A quiet ACE year still produces paid losses. Water, freeze, and severe convective storm do not wait for the basin to catch up. The cedent that treats September silence as a green light to cut CAT load mid-term is trading one verifying El Nino season against a catalog that still contains the last decade of inland flood and derecho.

What does not automatically change on the policy

- **Percentage hurricane and named-storm deductibles** in Florida, Texas, and the Gulf. Those attach to form definitions—named storm, hurricane, or windstorm—not to whether ACE matched the long-term mean in a given year. A quiet season does not commute them. See how limits and deductibles interact at claim time.
- **CAT load** in homeowners and commercial property rate. Load is an expected-loss construct across a catalog of seasons, not a rebate for the season in progress.
- **Reinsurance recoveries.** If the occurrence never crosses the attachment, the cedent keeps the attritional water and convective-storm burn. That is a working-layer year, not a free year.
- **Builders risk water deductibles and sublimits.** Market notes through 2026 still treat water as a leading builders-risk driver with rising deductibles. A missing Atlantic hurricane does not refill those sublimits.

The insured who wants a mid-term cut in CAT load is asking the carrier to treat one verifying El Nino season as new climatology. Underwriting will not do that in October with two months of the official season still open.

Wind deductible mechanics owners still face

Percentage triggers versus flat dollar deductibles

Coastal homeowners and many commercial property forms separate “all other perils” deductibles from hurricane, named-storm, or windstorm deductibles stated as a percentage of Coverage A or total insured value. Two percent on a \$1 million building is \$20,000 before the carrier pays the first dollar of wind or named-storm damage—often in addition to cosmetic matching disputes and non-recoverable depreciation on older roofs.

Definitions matter more than the Saffir-Simpson label

Edouard- and Fay-class events can sit inside “named storm” deductibles even when the basin has zero hurricanes and historically low ACE. Read the form before promising the insured that “there was no hurricane, so the flat deductible applies.” That distinction is core property insurance literacy, not meteorology.

Claims-desk implications while ACE stays light

Independent adjuster surge never starts. Staff adjusters stay on ordinary water and wind files longer. Documentation standards rise because nobody is clearing hundreds of CAT files a week. Higher deductibles and non-recoverable depreciation—already cited in 2026 restoration-market commentary as suppressors of claim filing—keep small tropical-rain jobs off the books even when a named storm exists.

For commercial accounts, the relevant question remains occurrence wording: named storm versus hurricane versus windstorm. Delay risk also flips. In a CAT year the complaint is cycle time. In a sub-11 ACE October the complaint is scrutiny. Expect more requests for cause-and-origin, more questions on matching, and fewer

“just pay the mitigation and close it” instructions. That is not hostility. That is a desk that is not drowning—see claims management discipline in a non-CAT year.

Model and treaty posture into 1/1

A season still hunting its first hurricane will show up in 2026 experience as a low-hurricane Atlantic year. It will not, by itself, cut 2027 cat-model AAL if the vendor catalog still contains the last decade of inland flood and severe convective storm. Brokers shopping 1/1 should separate three conversations: Atlantic hurricane ACE, U.S. severe convective and flood, and Pacific activity. The East Pacific has run active—major hurricanes including Polo and Odalys in late September—while Atlantic ACE stayed in single digits; Pacific energy does not offset Atlantic accumulation maps in a coastal renewal.

El Nino–driven shear and dry air explain the Atlantic shortfall; residual Gulf and Caribbean paths still belong in the accumulation map through November 30. If you are renewing a coastal schedule in October, do not bid the account as if 2026 is the new mean. Bid it as a year that verified a weak Atlantic pattern so far, with an open season and an unchanged form. Put the ACE chart in the submission. Do not put a “no CAT this year” sentence in the binder.

Portfolio and PML work still keys off long catalogs, not one El Nino season. So what: ACE 10.8 through October 1 is a working-layer year so far, not a reason to delete percentage wind deductibles or to treat CAT load as optional. Bind the form that is on the page. Revisit load at renewal with a full season, not with a September headline.

Frequently asked questions

Does a quiet hurricane season reduce my percentage wind deductible?

No. Percentage hurricane and named-storm deductibles are contract terms tied to how the policy defines the peril, not to whether ACE matched the long-term average in a given year. Carriers do not mid-term amend those triggers because the Atlantic had zero hurricanes through October.

What is ACE and why do insurers use it instead of storm counts?

Accumulated cyclone energy sums the intensity and duration of tropical systems at tropical-storm strength or higher. It tracks catastrophe exposure better than a count of weak, short-lived names. Rating, modeling, and treaty conversations use ACE and related loss metrics—not headline storm totals.

Can a tropical storm trigger a hurricane deductible?

Often no—but many coastal forms use broader “named storm” or “windstorm” deductibles that apply to tropical storms even when NHC never classifies a hurricane. The policy definition controls, not the Saffir-Simpson label on the evening news.

Will my carrier cut catastrophe load during the policy term if ACE stays low?

Generally not. CAT load reflects expected losses across many seasons in the rate filing, not a rebate for the current year. Mid-term requests to strip CAT load treat one quiet season as new science; underwriters typically defer that to renewal with full-year data.

Should I expect easier claim handling when there is no hurricane landfall?

Expect more scrutiny, not less. Without CAT surge, adjusters have capacity to dig into cause-and-origin, matching, and deductible math. Small tropical-rain claims may never meet the dollar threshold after a percentage wind or named-storm deductible—especially with non-recoverable depreciation on older roofs.

How should I renew a coastal property in October 2026?

Renew on the filed form and deductible structure, with two months of official season remaining. Submit ACE and outlook context, but do not underbid wind exposure because the Atlantic has not yet produced a hurricane. Separate Atlantic ACE from Pacific activity and from inland convective and flood loss trends in your submission narrative.

Cat Modeling & Property Risk Pulse: Active Tropical Storm Tracks, Surge Exposure & Commercial P&C Retentions — Tuesday, August 25, 2026

Source: <https://riskcoveragehub.com/cat-modeling-property-risk-pulse-active-tropical-storm-tracks-surge-exposure-commercial-pc-retentions-tuesday-august-25-2026/> Published: August 25, 2026

Updated October 6, 2026.

Direct Answer: This pulse records the NOAA NHC basin on Tuesday, August 25, 2026: twelve Atlantic and Eastern Pacific systems or investigative areas, weakening Tropical Storm Iselle west of Baja, and Atlantic AL96 at a 70% seven-day formation chance, while carriers still retained SCS-driven mid-layers and coastal named-storm deductibles near 2%–5% of TIV. In October 2026, use it as a dated retention and surge-context archive—then verify live NHC data and SOV-backed limits before placement.

Catastrophe Risk Briefing for Corporate Risk Managers, Commercial Underwriters, and Reinsurance Brokers: NOAA NHC basin developments, SCS loss accumulation, and commercial P&C treaty retention dynamics as of August 25, 2026.

Executive Catastrophe Exposure & Basin Activity Overview

On August 25, 2026, commercial property sat in peak convective and tropical season. P&C underwriters enforced tight attachment discipline and larger net retentions on secondary perils. NHC reported **12 active tropical systems and investigative areas** across the Atlantic and Eastern Pacific. SCS hail and straight-line wind remained the main mid-layer frequency driver—not a single U.S. landfalling major on that evening’s advisories. Model and accumulation context: catastrophe modeling guide and catastrophe portfolio management.

Tropical Cyclones & Atmospheric Inundation Tracking (NOAA NHC — August 25, 2026)

Outlooks from the evening of August 25, 2026 (archived pulse; live cones change continuously):

- **ATLANTIC (ABNT20 KNHC 252332, 800 PM EDT Tue Aug 25 2026): AL96** south-southwest of Cabo Verde—displaced convection, gradual development possible; tropical depression likely within a day or two, moving west to west-northwest; 48-hour and 7-day formation high, 70 percent; less favorable by the weekend. **AL95** northeast of Bermuda—elongated, hostile environment; 48-hour and 7-day formation low, 20 percent. Forecaster Papin.

- **EASTERN NORTH PACIFIC (ABPZ20 KNHC 252324, 500 PM PDT Tue Aug 25 2026):** Advisories on Tropical Storm Lala and Tropical Storm Iselle; final advisory on Remnants of Julio. Central East Pacific wave south of Mexico—depression expected in 2–3 days, 48-hour medium 60%, 7-day high 90%. Western East Pacific area east-southeast of Hawaii—48-hour medium 40%, 7-day high 70%. Forecasters Cangialosi/Katz/Evans.
- **Tropical Storm Iselle (EP09, Advisory 10, 200 PM PDT Tue Aug 25 2026):** Near 21.0N 121.8W (~775 mi west of Baja tip); 60 mph max winds; west-northwest 9 mph; no coastal watches; weakening to post-tropical Wednesday; Baja west-coast surf/rip risk; pressure 994 mb. Graphic timestamp: Tue, 25 Aug 2026 21:26:44 GMT.
- **Remnants of Julio (EP10, final Advisory 5):** Open wave near 21.0N 119.5W, north 25 mph, merging into Iselle’s circulation; last NHC advisory.

Live advisories: NOAA National Hurricane Center.

Surge & coastal exposure read-through (August 25 lens)

Iselle and Julio meant offshore wind and swell for Baja-facing coasts, not a U.S. Gulf or Atlantic surge event that date. AL96’s westward track still forced long-lead wind and surge reviews for Caribbean and Gulf-facing portfolios. Named-storm deductibles and flood sublimits follow policy structure—see Coverage A–D structure.

Cat Peril Loss Matrix & Commercial Underwriting Indicators (Late August 2026)

Peril Class	Primary Cat Models	Conditions as of Aug 25, 2026	Underwriting & Structuring Focus
Named Tropical Cyclone / Surge	Moody’s RMS, Verisk Touchstone	2%–5% named-storm TIV deductibles in coastal zones	Roof membrane age (<10 yrs) and flood gate certifications.
Severe Convective Storm (Hail/Wind)	Verisk, RMS, Swiss Re CatNet	Heavy mid-layer aggregation, Central and South US	Cosmetic hail exclusions, ACV roofs—see ACV, RCV, and agreed value.
Commercial Wildfire / WUI	CoreLogic Wildfire Risk, RMS HD	Capacity pressure into E&S surplus lines	Parametric wildfire liquidity on thermal triggers.
Inland Flash Flood / Atmospheric River	KatRisk, JBA Flood, FEMA NFIP	Loss frequency outside traditional 100-year SFHAs	Switchgear elevation audits and CBI sublimits.

Reinsurance Treaty Attachment & Capacity Dynamics (August 25, 2026)

Reinsurance capital had stabilized after multi-year property cat rate recalibrations, but attachments stayed elevated. Primaries retained larger net secondary-peril layers instead of ceding that volatility—see the reinsurance treaty guide. Insureds faced stricter co-insurance, water-damage sublimits, and SOV scrutiny.

Corporate Risk Manager & Insured Tactical Action Items (From This Pulse)

- **SOV precision:** Reconcile replacement cost valuations before co-insurance disputes.
- **Pre-loss MSAs:** Signed disaster response contracts for stacked SCS or tropical events.
- **Parametric liquidity:** Index wind triggers for deductible gaps—parametric and index-based risk transfer.

- **Envelope hardening:** RTUs, parapets, drainage ahead of convective and tropical squalls.

Where this stands now (October 2026)

Retrospective only—not live weather. August 25 showed a busy basin (EP storms and waves, AL96 at high formation odds, AL95 decaying north of Bermuda) while binders still centered SCS frequency and retained mid-layers. Iselle weakened offshore as NHC forecast; Julio was post-tropical. Confirm AL96’s later history in NOAA’s archive, not from this page.

By October 2026, readers should have refreshed SOV and valuations, checked surge/flood sublimits against gate certifications, and reconciled cat outputs with current pricing assumptions. Carry forward the discipline: archive the dated pulse, then update trajectories, treaties, and limits on live NHC and broker data.

Frequently asked questions

What was the commercial retention takeaway from the August 25, 2026 NHC outlook?

Carriers watched a full basin—AL96 at 70% seven-day formation, active Eastern Pacific storms, no U.S. landfalling major in that evening’s advisory set—while retaining SCS mid-layers. Coastal named-storm deductibles at 2%–5% TIV still framed surge and wind attachment even when headline storms stayed offshore.

Why did SCS losses still dominate the mid-layer story that week?

Hail and straight-line wind across the Central and South US drove heavy commercial aggregation that models capture unevenly by location. Underwriters pushed cosmetic hail exclusions and ACV roof schedules—where valuation method sets net recovery.

Did Tropical Storm Iselle on August 25, 2026 change U.S. Gulf or Atlantic surge planning?

No for U.S. Atlantic and Gulf coasts that date: Iselle was west of Baja, no watches, forecast post-tropical, with Baja swell/rip hazards. U.S. surge planning leaned on long-lead Atlantic waves such as AL96 more than Iselle’s track.

How should reinsurance brokers read the treaty paragraph in this pulse today?

As a late-August 2026 snapshot: stabilized reinsurance capital, elevated attachments, primaries holding net secondary-peril risk. October placements need current treaty terms—not this dated retention line alone.

Can I rely on this page for live storm coordinates?

No. Positions and outlook text are frozen to August 25, 2026. Use nhc.noaa.gov for active coordinates, cones, and warnings.

What is the first October 2026 action after revisiting this pulse?

Pull current NHC outlooks, then reconcile SOV, named-storm and flood deductibles, and parametric or treaty gaps against today’s bound program with live broker and engineer input.

Official Citations & Authoritative Sources

- **NOAA National Hurricane Center (NHC):** nhc.noaa.gov
- **NOAA Storm Prediction Center (SPC):** spc.noaa.gov

- **Moody's RMS Catastrophe Modeling:** rms.com
- **Verisk Extreme Event Solutions:** verisk.com

Underwriting & Risk Analysis Disclaimer: For commercial risk managers, brokers, and property owners—situational awareness only, not underwriting guidance or a binder. Verify weather with NOAA/NHC; verify terms with licensed brokers and engineers.

Related on this hub

Also on Risk Coverage Hub: catastrophe modeling complete guide, reinsurance treaty guide, and climate risk pricing and catastrophe model updates.

Scope 3 Supply Chain Data and the Insurance Claims Ecosystem: What the Restoration Carbon Protocol Means for Carriers

Source: <https://riskcoveragehub.com/restoration-carbon-protocol-insurance-carriers-scope-3-claims/>

Published: April 11, 2026

Updated October 6, 2026.

Direct answer: Paid property claims pull restoration contractors into a carrier's Scope 3 supply chain, but emissions data usually never reaches the claim file unless the broker and claims operation require it. The Restoration Carbon Protocol (RCP) is an open, per-job standard that lets carriers ingest contractor carbon data in a consistent, audit-ready format alongside ordinary claim documentation.

By October 2026, U.S. property carriers still face overlapping climate disclosure pressure on the claims side of the house—not on catastrophe models or treaty renewals, but on whether they can describe and evidence emissions from purchased restoration services. The NAIC Climate Risk Disclosure Survey for reporting year 2025 was due August 31, 2026 from insurers licensed in participating states with more than \$100 million in nationwide direct written premium; responses follow the TCFD-aligned questionnaire the NAIC adopted in 2022 and are collected through the California Department of Insurance multi-state portal. Federal SEC climate rulemaking remained in flux after the May 2026 proposal to rescind the 2024 rules, which does not cancel state insurance surveys or carrier sustainability reports.

Internationally, more than 40 jurisdictions have adopted or are introducing ISSB-based sustainability disclosure, with IFRS S2 climate requirements phasing in on jurisdiction-specific dates (many first reports landing in 2025–2027). EU-exposed carrier groups face a CSRD landscape reshaped by the 2025 “stop-the-clock” postponement and the March 2026 scope narrowing in Directive (EU) 2026/470: large public-interest insurers already in the first CSRD wave continue reporting, while many undertakings previously scheduled for 2025–2026 reporting years were postponed or lifted above a 1,000-employee threshold effective for financial years starting on or after 1 January 2027. Wherever those frameworks apply, Scope 3 from purchased services remains a recurring gap—and restoration work on indemnity claims is a concrete, high-volume purchased-services line item.

Scope 3 data from a restorer does not land on a carrier desk by itself. It has to survive the producer who still places the contractor—the seat John Hubbard occupies at USI. If that person cannot explain the number, the claim file never carries it. Downstream, claims teams need a place in the claims management workflow for optional sustainability fields without breaking property claim filing and documentation discipline from first

notice through settlement.

Where contractor emissions sit in a carrier's Scope 3 inventory

For insurance carriers, restoration contractor emissions are Scope 3 Category 1 (purchased goods and services) under the GHG Protocol Corporate Standard. Carriers with Science Based Targets initiative commitments typically include Category 1 inside the minimum Scope 3 coverage boundary. Carriers reporting under CSRD (where still in scope), ISSB-aligned regimes, or the NAIC survey's metrics narrative still need a defensible story for purchased services tied to real operations—not only investments and owned facilities.

The current state in most carrier sustainability reports: restoration and broader claims-vendor emissions are either excluded with an explicit data-limitation note, or estimated with spend-based factors (for example EPA USEEIO-style construction and services proxies). Both approaches are widely flagged as low data-quality tiers. As limited and reasonable assurance expectations tighten around GHG inventories, spend proxies for field services that vary by drying strategy, equipment mix, and disposal path draw scrutiny—especially when institutional investors compare carriers on Scope 3 methodology notes.

What the claim ecosystem must capture

A usable supply-chain record for one loss needs more than a portfolio emissions ratio. At minimum, carrier and broker stakeholders ask whether the file can show:

- Which restoration vendor performed the work and under which managed-repair or preferred-vendor program
- Job-level activity data (equipment run time, fuel use where relevant, waste streams, material choices) rather than only invoice total
- Methodology, emission factors, and data quality tier documented for assurance reviewers
- Alignment with how the carrier already handles regulatory compliance evidence—retention, audit trail, and named ownership

That is a claims-data problem first and a sustainability-reporting problem second. It intersects with broader regulatory compliance hygiene: if nobody owns the field, the disclosure stays a footnote.

What the Restoration Carbon Protocol provides

The Restoration Carbon Protocol (RCP) v1.0, published at tygartmedia.com/rcp, is a free, open-source standard for per-job Scope 3 emissions calculation from restoration contractor work. It is the primary public specification for RCP's methodology, schema, and reporting artifacts. For insurance carriers evaluating vendor standards, RCP is not a replacement for policy terms or claim payment rules; it is a data contract for carbon attributable to a single restoration job.

- **Per-claim carbon data:** Each RCP Job Carbon Report covers one restoration job—actual job-level data rather than a portfolio spend roll-up.
- **Machine-readable format:** The RCP-JCR-1.0 JSON schema is intended for intake into ESG data management platforms, vendor scorecards, and sustainability reporting workflows—similar in spirit to how carriers are hardening other third-party data feeds under insurance regulatory technology governance, but aimed at post-loss services.

- **Audit-ready methodology:** RCP records document calculation methodology, emission factor sources, and data quality tier in a structure aligned with GHG Protocol verification practice for Scope 3 activity data.
- **Carbon avoidance documentation:** The framework also quantifies avoided emissions—for example when a contractor dries in place rather than demolishes, or diverts debris from landfill—supporting climate action narrative and, separately, factual loss-mitigation documentation in the claim file.

How carriers and brokers wire RCP into the claims supply chain

Implementation is vendor management and claims operations, not a separate ESG side project. Practical insertion points include:

- **Preferred vendor and managed repair criteria:** Require an RCP Job Carbon Report (or equivalent mapped fields) before final invoice approval or quality review on qualifying property claims.
- **TPA and desk-adjusting instructions:** Add optional sustainability fields to FNOL and repair monitoring checklists so adjusters request reports when mitigation scope is known.
- **Broker-led placement:** Retail and wholesale producers attach RCP expectations to contractor panels they certify for carrier programs—so the number is explained before the carrier’s sustainability team asks for it.
- **ESG and finance intake:** Map RCP-JCR JSON into the same inventory pipeline used for other Category 1 services, with job IDs cross-walked to claim numbers for assurance sampling.

Carriers already running NAIC survey narratives should treat contractor carbon the way they treat other emerging metrics: name the owner, describe the boundary, and state whether data are measured, estimated, or excluded. The August 2026 reporting year 2025 cycle is closed for late filers; reporting year 2026 planning should assume the TCFD-aligned survey continues regardless of federal securities rule outcomes.

The carrier network effect in restoration

The practical leverage of an open protocol in insurance is supply-chain breadth. When a carrier names RCP—or a materially equivalent open schema—as a preferred-vendor deliverable, every contractor in that panel faces the same implementation target. Data quality improves when adjusters and vendors repeat the same job report on every qualifying loss instead of one-off spreadsheets at year-end.

EcoClaim, a Canadian company, has built a proprietary, carrier-facing version of the same concept for Canadian P&C markets: closed platform, subscription model, carrier-specific data stores. RCP presents an open-standard path: any carrier can reference the public specification, any contractor can implement without a platform fee, and multiple carriers can ingest the same JSON shape into different disclosure programs. Partnership and governance questions for industry adoption are directed through the contact listed on tygartmedia.com/rcp (rcp@tygartmedia.com).

Scope 3 from restoration is not abstract climate risk—it is paid claim activity flowing through licensed vendors, TPAs, and broker panels. Carriers that standardize per-job data at that choke point improve both disclosure defensibility and ordinary governance over third-party records. Brokers who understand the metric keep it on the file; carriers who specify the format receive supply-chain data that can survive assurance review instead of dying in a spend-based proxy footnote.

Related on Risk Coverage Hub: Home.

Frequently asked questions

Does the NAIC Climate Risk Disclosure Survey require restoration contractors to submit RCP reports?

No. The NAIC survey is a TCFD-aligned insurer questionnaire; it does not mandate a contractor file format. Carriers describe metrics, targets, and data limitations in their own responses. RCP is one voluntary method to generate job-level Category 1 data carriers may reference when they explain purchased-services emissions in survey narratives and sustainability reports. Which Scope 3 category covers restoration vendors on paid property claims?

Restoration and mitigation invoices are generally Scope 3 Category 1 (purchased goods and services) for the carrier paying the vendor. They are not Scope 1 or 2 unless the carrier owns the equipment and employs the crews directly, which is uncommon on standard indemnity claims. How is RCP different from spend-based EPA or USEEIO estimates?

Spend-based models multiply invoice dollars by economy-wide emission factors. RCP is built for per-job activity data from restoration work—equipment, energy, materials, and disposal paths—so the inventory can move from generic spend tiers toward supplier-specific activity data where contractors participate. Where should a carrier require an RCP Job Carbon Report in the claim lifecycle?

Most carriers attach it to vendor compliance at invoice or completion: preferred vendor programs, managed repair guidelines, or TPA standard operating procedures after scope is fixed. The report should reference the claim or work order identifier used in core claims systems so assurance teams can sample jobs. Do CSRD and ISSB S2 rules force U.S. carriers to collect contractor carbon on domestic claims?

Not automatically. CSRD applies to in-scope EU entities and certain group parents; ISSB adoption depends on each jurisdiction's rules and effective dates. U.S. carriers with EU subsidiaries or ISSB-aligned listing requirements may still need Category 1 data that includes purchased restoration services. Domestic NAIC-participating insurers face separate TCFD-aligned survey expectations. The claims-supply-chain data problem is shared; the legal hook varies by entity and jurisdiction. Can brokers block Scope 3 data even when a carrier wants it?

Yes, if the producer does not contract for it or cannot explain it to the contractor panel. Carriers set policy; brokers and TPAs execute placement and instructions. Without broker-facing language and panel rules, contractor reports stall before they reach the carrier's ESG intake.

Climate Risk and Insurance Pricing in 2026: How Physical Hazards Are Repricing Every Line of Coverage

Source: <https://riskcoveragehub.com/climate-risk-insurance-pricing-2026-physical-hazards/> Published: April 3, 2026

Updated October 6, 2026.

In 2026, property and casualty carriers price physical climate risk into premiums, limits, and geographic appetite using forward-looking catastrophe models—not decades of loss history alone. Coastal, wildfire, flood, and severe-convective exposures face double-digit rate increases, tighter sub-limits, and selective non-renewals while reinsurance and catastrophe bonds carry more peak risk at higher cost. Treat quotes and declinations as market signals and pair traditional coverage with mitigation, alternative transfer, and documented retention.

Carriers now embed climate-adjusted loss projections in pricing, limits, and geographic appetite across commercial property, homeowners, business interruption, workers compensation, and auto. That is not a slow drift: non-renewals, moratoriums, and exits in high-hazard states compress options while reinsurance and excess markets stay firm.

Catastrophe modeling in 2026: backward-looking to forward-looking

For decades, carriers relied on CAT models from firms such as Moody's RMS, Verisk AIR, and CoreLogic EQECAT. Those tools fit historical loss catalogs, generated synthetic event sets, and worked when climate was relatively stable. That assumption no longer holds. Flood frequency, hurricane intensity, wildfire geography, hail and tornado patterns, and temperature extremes are shifting faster than a static history can represent.

Vendors blend history with downscaled climate projections under multiple emissions paths. Updated catalogs show higher expected losses in exposed counties, feeding rate filings and portfolio decisions. See *Climate Risk Pricing and Catastrophe Model Updates (2026)* and *Catastrophe Modeling: The Complete Guide*.

Synthetic event catalog: A statistically representative sample of thousands of hypothetical future loss events (earthquakes, hurricanes, floods, and similar perils) generated by catastrophe models. Each event has location, severity, and loss estimates so actuaries can estimate portfolio loss distributions.

Premium increases outpacing loss experience

In competitive, stable markets, premiums track loss experience. In a hard market, rates rise faster than paid losses because capacity is scarce and forward risk dominates backward data. That imbalance is acute in 2026 across coastal counties, wildland-urban interface zones, and flood-prone parcels: property premiums in many of those areas have moved up 30–50% or more over roughly three years while recent loss ratios alone do not explain the full increase on a historical basis.

Carriers are charging for climate-adjusted expected loss, not just the last three storm seasons. Homeowners and commercial owners face an affordability squeeze when premium growth outpaces asset values—what regulators and legislators increasingly call an insurance affordability crisis. Commercial operators see the same pressure on net operating income and lender covenants. The cycle mechanics—capacity, attachment points, and filing politics—are covered in *Hard Market vs Soft Market in Insurance (2026)*.

Coverage narrowing: what gets excluded

Tightening appetite shows up as geographic exit, moratoriums on new business, and selective non-renewals. California, Florida, and Louisiana remain the headline states, but the pattern is peril-specific everywhere: water damage and flood exclusions even outside FEMA-mapped zones, wind restrictions on coastal forms, wildfire moratoriums in the West, and stricter enforcement of policy conditions after events.

Sub-limits shrink the effective cover. Earthquake caps on large schedules, water damage caps unrelated to building values, and business income waiting periods or limited indemnity periods leave meaningful retention on the balance sheet. Policyholders reviewing ACV, RCV, and agreed-value structures should read endorsements line by line; *Property Insurance: The Complete Professional Guide (2026)* walks through how carriers document those choices in 2026 filings.

Sub-limit: A maximum payout for a listed peril inside a broader policy. A \$10 million building schedule with a \$100,000 earthquake sub-limit leaves everything above that cap uninsured for shake damage.

Reinsurance and catastrophe bonds: cost at the top of the tower

Reinsurance still absorbs peak catastrophe layers for primary carriers. After heavy 2023–2025 global catastrophe years, reinsurers pushed renewal pricing and tightened terms into 2026. Higher reinsurance spend flows straight to primary rates or to reduced limits on cat-exposed books.

Alternative capital is not a sideshow. Artemis tracked record first-half 2026 catastrophe bond issuance of roughly \$18 billion across 83 Rule 144A and private transactions, with outstanding market size near \$65.6 billion at mid-year—both benchmarks for the sector. Sponsors use those bonds to collateralize reinsurance layers; when cat bonds price tighter or widen, primary carriers feel it in their tower economics. Hybrid programs—traditional reinsurance plus index triggers—are common on large commercial schedules; see Parametric Insurance and Index-Based Risk Transfer for how buyers pair indemnity with collateralized layers.

Parametric insurance: growth and limits

Parametric coverage pays when an objective index or trigger is met—rainfall depth, wind speed at a reference station, earthquake magnitude band—without waiting for a full indemnity adjustment. Settlement is faster and moral hazard is lower, but basis risk is real: the index can pay when your site is unscathed, or fail to pay when localized damage is severe.

Parametrics work best where peril and index correlate tightly (many flood and wind designs). Wildfire and hail are harder at portfolio scale because fuel, topography, and swath width vary block by block. In 2026, many buyers use parametric layers for immediate liquidity after an event and indemnity forms for rebuild cost—especially where NFIP or private flood markets leave gaps on commercial schedules.

Basis risk: The mismatch between an index payout and actual loss. A triggered \$500,000 parametric flood payment with \$200,000 of real damage leaves surplus cash; damage of \$800,000 with no trigger leaves a hole.

Social inflation stacked on physical climate trends

Construction and medical inflation still inflate claims. Social inflation—larger verdicts and settlements in liability lines—adds another layer. Climate-driven events that implicate utilities, municipalities, or employers can produce outsized liability stacks on top of property damage. Carriers modeling both frequency and severity drift are repricing general liability, property, workers compensation in outdoor and heat-exposed classes, and homeowners in wildfire corridors accordingly.

Market segmentation and the uninsurable gap

Low-hazard locations still see competitive markets. High-hazard locations face FAIR-plan or insurer-of-last-resort programs, captives, or bare retention. Lenders, investors, and bond covenants treat evidence of insurance as a gate; when standard markets walk away, financing and valuation follow.

State regulators continue to press carriers on rate adequacy, market conduct, and climate disclosure expectations even as federal securities climate rules fluctuate. The NAIC climate-risk survey cycle remains a practical compliance checkpoint for insurers regardless of SEC rule timing. Filing disputes and consumer

complaints still route through state departments of insurance when coverage is reduced or denied.

What risk managers should do now

Mitigate physically. Defensible space, ignition-resistant construction, elevation and dry floodproofing, roof and envelope upgrades, and heat protocols reduce modeled loss and sometimes earn credit in underwriting files.

Document retention. Know sub-limits, coinsurance, waiting periods, and ordinance-or-law gaps before a loss. Align BI worksheets with actual payroll and dependency maps.

Transfer creatively. Layer parametrics, captives, or cat bonds where traditional capacity stops; score locations with the same rigor carriers use—Risk Scoring and Insurance Underwriting describes how underwriters translate hazard scores into declinations.

Read market signals. Clustered non-renewals are a hazard signal, not a marketing glitch. Where disclosure rules apply, document insurance friction as evidence of material physical risk.

Conclusion

Insurance pricing in 2026 reflects a reset in how the industry underwrites physical climate risk. Forward-looking catastrophe models, expensive reinsurance, and record cat-bond issuance are pulling the same direction: higher cost, narrower forms, and sharper geography. Organizations that treat insurance as one layer in a broader resilience plan—mitigation, retention, alternative transfer, and regulatory engagement—are aligned with where the market is already going. Cheap coverage in high-hazard places is not coming back on a historical curve alone.

Frequently asked questions

How have catastrophe models changed in 2026 to account for climate risk?

Traditional CAT models relied on decades of historical loss data to build synthetic event catalogs. Updated 2026 models blend that history with climate projections from downscaled global circulation models, adjusting frequency and severity for perils such as flood, hurricane wind, wildfire, and severe convective storm. The result is higher modeled loss in many exposed geographies, which carriers use in rate filings, reinsurance purchases, and geographic cutoffs.

Why are insurance premiums rising faster than recent loss experience alone would suggest?

Carriers price for expected future loss under climate-adjusted scenarios, not only the last few years of paid claims. Where capacity is limited, competition does not force rates down to historical experience. In high-hazard counties, double-digit cumulative increases over several years often reflect forward models plus higher reinsurance and cat-bond costs passed through to policyholders.

What did the catastrophe bond market do in the first half of 2026?

Artemis reported record first-half 2026 catastrophe bond issuance of roughly \$18 billion across 83 transactions, with outstanding market size near \$65.6 billion at mid-year. That activity shows institutional capital still absorbing peak peril layers, but pricing and collateral terms still flow through to primary insurers and, ultimately, commercial and personal lines pricing.

How does parametric insurance differ from traditional indemnity coverage?

Parametric policies pay when a predefined index or trigger is met, often within days, without a full loss adjustment. Traditional indemnity pays proved physical loss subject to policy terms. Parametrics can fund immediate cash flow but carry basis risk; most large buyers in 2026 use them as a liquidity layer above or beside indemnity programs.

What should property owners do when standard markets non-renew or cap coverage?

Start with verifiable mitigation and complete insurance data (valuations, COPE details, loss history). Shop surplus lines and specialty flood or wind markets where admitted carriers exit. Consider captives, parametrics, or higher deductibles with funded reserves. Engage state regulators when non-renewal patterns look discriminatory by peril rather than by risk quality, and document how uninsurable exposure affects loans, leases, and disclosure obligations.

Climate Risk Pricing and Catastrophe Model Updates: How Insurers Quantify Escalating Natural Disaster Losses

Source: <https://riskcoveragehub.com/climate-risk-pricing-catastrophe-model-updates-2026/> Published: April 1, 2026

Updated October 6, 2026.

Property insurers are repricing with forward-looking catastrophe models as secondary perils—wildfire, severe convective storms, and flood—drive most loss. Reinsurance softened at January 2026 renewals while retail rates still climb; catastrophe bonds set an H1 record near \$18 billion across 83 deals (Artemis).

Catastrophe modeling and climate risk pricing

Cat models simulate hazards and map them to insured loss; climate risk pricing feeds those outputs—average annual loss, tails, regulatory scenarios—into premium, deductibles, and capacity. In 2026 the story is secondary-peril persistence and models that no longer treat history as stationary.

Loss and pricing in 2026

Swiss Re Institute estimates 2025 global insured natural-catastrophe loss at about USD 107 billion (190 events)—lowest in five years but still elevated. Secondary perils hit a record 92% share (USD 98.3 billion); wildfire, severe convective storms, and flood alone were 88%. January 2025 Los Angeles-area wildfires drove roughly USD 40 billion insured; global severe convective storm loss stayed near USD 51 billion. Swiss Re's long-run trend still implies ~5–7% real annual growth; a return to trend could land near USD 148 billion in a normal year.

Primary homeowners rates keep rising despite softer reinsurance—often cited near 4% in 2026 nationally (~USD 3,057), with sharper moves in California and several central states and Florida still highest. See climate risk and insurance pricing in 2026.

California wildfire filings

California's Sustainable Insurance Strategy lets carriers use Department-reviewed forward-looking wildfire cat models and net reinsurance in rate filings, with write commitments in wildfire-distressed ZIP codes to shrink FAIR Plan reliance. Verisk, Karen Clark & Company, and Moody's RMS completed review. Trade-off: higher modeled premiums in fire zones versus better voluntary-market availability when carriers can price to survive.

Models, reinsurance, and cat bonds

Vendors (Moody's RMS, Verisk, CoreLogic/Cotality, KCC, Reask, Jupiter) rolled climate-conditioned and secondary-peril updates in 2025–2026—higher forward hail/wind/wildfire loss costs in many portfolios. Insurers blend vendor views with own-view-of-risk; NAIC property RBC filers report climate-conditioned hurricane and wildfire impacts using approved catalogs. Start with catastrophe modeling: the complete guide (2026).

Howden Re put risk-adjusted global property-cat reinsurance rates-on-line down ~14.7% at January 1, 2026—the steepest drop since 2014—as capital stayed strong and 2025 lacked a major U.S. hurricane landfall. Retail rates lag filings and modeled loss inflation; see hard vs soft market (2026).

Artemis tracked ~USD 18 billion of Rule 144A and private cat-bond issuance across 83 transactions in H1 2026—a first-half record—as sponsors added wildfire, quake, and secondary-peril covers. Triggers and structures tie to parametric insurance and catastrophe bonds; align limits with modeled PML via catastrophe portfolio management.

Regulation and disclosure

The NAIC Climate Risk Disclosure Survey for reporting year 2025 is due August 31, 2026 for large writers in participating states, including cat-model use for climate risks. Federal SEC climate rules may shift, but state hooks remain—see SEC rescission and the NAIC survey timeline. NFIP Risk Rating 2.0 continues to phase through federal flood policies; private flood cat modules still run ahead of political caps on NFIP pricing.

Action checklist

- Re-rate locations on forward peril views, not 2010s ZIP memory.
- Split wind/hail/flood sub-limits; stress secondary perils apart from hurricane.
- Document mitigation for underwriting and claims.
- Compare 1.1 reinsurance savings with cat-bond/parametric capacity where basis risk is acceptable.

Frequently Asked Questions

How large were 2025 insured catastrophe losses?

Swiss Re Institute estimates about USD 107 billion globally. Secondary perils—wildfire, severe convective storms, and flood—were 92% of the total, including roughly USD 40 billion from January 2025 Los Angeles-area wildfires and about USD 51 billion from severe convective storms.

Why did reinsurance fall in January 2026 while homeowners rates rose?

Howden Re reported risk-adjusted global property-cat rates-on-line down about 14.7% at January 1, 2026 renewals, helped by strong capital and no major U.S. hurricane landfall in 2025. Primary carriers still face

higher modeled loss costs, rebuild inflation, and filing lag, so retail property rates kept climbing in many states.

What happened in cat bonds in H1 2026?

Artemis recorded about USD 18 billion of issuance across 83 Rule 144A and private cat-bond transactions in the first half of 2026—a first-half record—as sponsors placed more U.S. wind, wildfire, earthquake, and secondary-peril risk in ILS form.

Can California use cat models in rate filings?

Under California’s Sustainable Insurance Strategy, insurers may use Department-reviewed forward-looking wildfire cat models and net reinsurance in filings, with commitments to write specified shares in wildfire-distressed ZIP codes. Verisk, Karen Clark & Company, and Moody’s RMS models completed review.

What are regulators asking about cat models in 2026?

The NAIC Climate Risk Disclosure Survey for reporting year 2025 asks large participating insurers how they identify climate-related risks and whether they use catastrophe modeling to manage them, due August 31, 2026. Property RBC filers also report climate-conditioned hurricane and wildfire modeled impacts.

Catastrophe Modeling: Expert Video Analysis

Source: <https://riskcoveragehub.com/video-catastrophe-modeling/> Published: March 18, 2026

Updated October 6, 2026.

Direct answer: This page centers on a walkthrough of natural catastrophe modeling for insurance and reinsurance audiences—how stochastic event catalogs, vulnerability curves, and financial engines roll up to metrics carriers use for pricing, accumulation control, and capital planning. Watch the embedded video first for the full narrative arc; the sections below summarize what it covers, what still holds in October 2026, and where to go deeper on Risk Coverage Hub without replacing the video.

Risk Coverage Hub › Video resource

Source video: *Natural Catastrophe Modeling in Insurance and Reinsurance* | NatCat Risk Management Explained (YouTube, November 17, 2024). Risk Coverage Hub added practitioner notes and October 2026 context around the embed.

Why start with this video

Cat models sit behind decisions most property owners never see directly: territory pricing, minimum wind or wildfire deductibles, aggregate caps, and whether a location clears underwriting. The walkthrough explains the discipline in plain insurance terms—what goes in, what comes out, and why two competent teams can still disagree on tail loss for the same address.

If you need a written baseline before or after you watch, pair the video with *Catastrophe Modeling: The Complete Professional Guide* (2026). That guide stays at framework level; this page keeps the moving picture up front.

What the video explains

The presentation follows the standard three-engine architecture commercial vendors use, without getting lost in vendor marketing slides.

Hazard and stochastic events

Hazard modules draw from history, physics, and statistical extrapolation to build large event catalogs—hurricane tracks, earthquake ruptures, severe convective footprints, and other perils depending on scope. Each simulated event carries a frequency and a geographic intensity field so every insured location receives an intensity measure for that scenario.

Vulnerability and exposure

Damage functions translate intensity at a site into physical loss given construction class, occupancy, height, year built, and modifiers such as roof type or anchorage. Exposure is the schedule of values and attributes feeding those functions. Sloppy geocoding or stale replacement costs move tail metrics more than most policyholders expect; models assume the SOV matches underwriting files.

Financial module and roll-up metrics

The financial layer applies limits, deductibles, coinsurance, sublimits, waiting periods, and reinsurance terms where modeled. Outputs include average annual loss (AAL), occurrence and aggregate exceedance probability curves, and tail summaries used in reinsurance submissions and internal capital discussions. Policy language still governs what pays on a real claim—modeled limits must match bound forms and endorsements, a topic we treat in Insurance Policy Coverage Analysis: ISO Forms, Endorsements, and Coverage Gaps.

Validation and model change

The video treats validation as ongoing hygiene: compare modeled distributions to credible historical benchmarks, stress recent major events, and rerun portfolios when vendors ship new catalogs or vulnerability suites. Model change is an operating rhythm, not a one-time IT project.

Key takeaways after you watch

- Cat models answer portfolio questions across thousands of synthetic years; they do not predict the date of the next landfall or rupture.
- Vendor platforms—chiefly Moody’s RMS and Verisk AIR in U.S. property—share architecture but diverge on catalogs, damage assumptions, and financial granularity. Material spread at the tail is normal; document which version powered your renewal analytics.
- Exposure quality beats brand debates. Correct coordinates, values, occupancy, and construction drive results more than small differences in event rates.
- Underwriting teams use modeled central tendency for rate adequacy and tail metrics for capacity and aggregation limits—often long after a quiet season ends. See Hard Market vs Soft Market in Insurance (2026) for how that shows up in coverage and pricing even when recent loss counts are low.
- Investors and sponsors price insurance-linked securities with the same tail logic; strong cat bond issuance in 2026 reflects demand for modeled risk transfer, not proof that physical losses accelerated overnight.

October 2026 context for model users

Vendor release cycles continue to emphasize expanded secondary-peril modules—severe convective storm, wildfire, inland flood where offered—and optional climate-conditioned hazard views for selected regions. Treat baseline versus forward-looking runs as different labels on the same portfolio; year-over-year comparisons fail when the catalog changes silently.

On the investor side, Artemis tracked roughly \$17.98 billion of new Rule 144A and private catastrophe bond issuance across 83 transactions in the first half of 2026—a record first half—with outstanding market size near \$65.6 billion at mid-year (Artemis.bm). Those deals lean on modeled exceedance probabilities and expected loss analytics at attachment points; model credibility at the tail still drives pricing.

On the public-sector education side—not a substitute for carrier analytics—the NAIC adopted a Catastrophe Modeling Primer in March 2025 to introduce probabilistic models to regulator audiences, and NAIC forums in 2026 discussed piloting near-real-time loss-estimation tools for state teams ahead of landfall or major convective outbreaks. That work sits in education and preparedness lanes; your renewal still rides on carrier vendor stacks and clean exposure files.

Flood remains an exposure footnote in many property cat runs: NFIP versus private flood coverage, elevation, and program rules determine whether flood loss appears in a given portfolio analytics pass. FEMA publishes NFIP technical references at [fema.gov/flood-insurance](https://www.fema.gov/flood-insurance); cat teams still need accurate location and limit data regardless of program.

For vendor-specific mechanics and metric definitions (AAL, OEP, AEP, TVaR), continue with Catastrophe Modeling: How RMS, AIR, and Verisk Quantify Catastrophe Risk.

How property owners and risk managers should use this page

Watch the embed once for orientation, then skim the takeaways before a renewal where your carrier asked for updated values, roof documentation, or wildfire mitigation proof. Ask your broker whether modeled concentration in your territory—not last year’s storm count—drove deductible or sublimit changes. Bring an accurate SOV; it is the fastest way to avoid conservative defaults when files are thin.

Frequently asked questions

What is catastrophe modeling in plain language?

Catastrophe modeling simulates rare, severe events across many synthetic years and translates them into insured loss distributions for a book of business. It combines hazard science, engineering judgment about damage, and insurance math about deductibles and limits.

Do I need to license RMS or AIR as a property owner?

Most owners rely on insurer or broker analytics rather than licensing a vendor platform. What you control is exposure quality: addresses, construction, occupancy, values, and business income assumptions. Poor inputs distort tail metrics more than choosing one vendor name over another.

How should I interpret average annual loss versus a 100-year modeled loss?

Average annual loss is the mean across the entire simulation and supports expected-loss thinking. A 100-year occurrence or aggregate point is a tail metric on an exceedance curve, used for capacity and reinsurance-style

decisions. Use both—central tendency alone hides severe years.

Why do wind or wildfire deductibles stay high after a quiet catastrophe season?

Models price long-horizon risk, not last season's event count. Carriers also watch modeled aggregation, reinsurance cost, and vendor model updates. A low-storm year does not automatically reset deductibles that were set using tail metrics.

What should I read next on Risk Coverage Hub after watching?

Start with the complete professional guide and the RMS, AIR, and Verisk article linked above, then read about hard versus soft market cycles if your renewal pricing feels disconnected from recent weather. Keep policy forms in view whenever modeled limits are discussed.

How do catastrophe bonds relate to the models in the video?

Cat bond investors price notes using modeled loss distributions and attachment probabilities for the covered peril or portfolio. Record first-half 2026 issuance reflects sponsor demand for that transfer mechanism; it does not mean every modeled peril suddenly worsened in one quarter.

Hurricane, Earthquake, and Wildfire: Catastrophe Peril Analysis for Insurance Professionals

Source: <https://riskcoveragehub.com/hurricane-earthquake-wildfire-catastrophe-peril-analysis/> Published: March 17, 2026

Updated October 6, 2026.

Direct answer: Hurricane, earthquake, and wildfire losses are driven by different physics and building vulnerabilities, so property programs must treat them as separate perils—not one generic “catastrophe” bucket. Wind and storm surge dominate hurricane claims; ground shaking and code era drive earthquake severity; ember exposure, fuel load, and defensible space drive wildfire survivability. Standard property forms often cover hurricane wind but exclude flood and earthquake, and wildfire underwriting now ties closely to mitigation and last-resort markets where admitted capacity is thin.

This peril-by-peril view is for owners, risk managers, contractors, and brokers aligning coverage, deductibles, and retrofit decisions with how each event damages buildings. For program context, see property insurance: the complete professional guide.

Why peril-specific analysis beats a single “cat” label

Portfolio stress tests and reinsurance towers matter, but at the building level the failure modes differ. A coastal frame house loses its roof deck in a hurricane; the same structure may ride out shaking if it is a ductile moment frame, yet burn from an ember two miles ahead of a flame front. Treating all three as interchangeable drives wrong deductibles, missing endorsements, and repair scopes that ignore code upgrades triggered by each peril.

Carriers still tighten terms in hard markets—higher wind and wildfire deductibles, tighter roof and age rules, and more scrutiny of reconstruction cost—without changing the underlying physics. Understanding each peril keeps you ahead of declinations and claim disputes before renewal.

Hurricane peril analysis

Physical characteristics and loss drivers

Hurricanes combine sustained wind, gusts, wind-driven rain, tornado embeds, and coastal storm surge. Onshore loss severity usually tracks maximum sustained wind, exposure to open fetch, roof cover type and attachment, opening protection (doors, windows, garage doors), and interior water entry after envelope failure. Surge and heavy rainfall produce flood and business-interruption losses even when wind speeds are moderate.

Secondary damage—mold, corrosion, and extra expense—often exceeds the initial wind line item if drying and power restoration lag. On large commercial schedules, contingent time element and civil authority clauses can dominate economic loss when evacuation orders persist.

Building code and construction vulnerabilities

Post-storm forensics repeatedly show roof-to-wall connections, gable end bracing, soffit integrity, and garage door pressure rating as decisive. Older coastal stock built before modern wind design codes suffers disproportionate uplift and water intrusion. Low-slope commercial roofs fail at seams and penetrations; masonry veneers crack when flexible diaphragms are absent.

Surge adds hydrodynamic load, debris impact, and foundation scour. Elevated structures on truly breakaway enclosures with elevated utilities perform better. Mixed-use buildings with rooftop equipment need anchor plans; unsecured HVAC units become projectiles.

Coverage and underwriting issues

Commercial property and homeowners policies typically cover wind from hurricanes unless restricted by endorsements, but **flood—including surge—is excluded** and must be addressed through the National Flood Insurance Program (NFIP) or private flood markets. FEMA fully implemented **Risk Rating 2.0** by April 1, 2023; premiums reflect individual flood frequency, distance to water, elevation, and replacement cost rather than broad zone averages alone. FEMA continues annual updates and publishes state transition profiles at fema.gov/flood-insurance/risk-rating. New business still rates on property-specific flood characteristics, not legacy zone tables.

Named-storm, hurricane, or wind/hail deductibles apply as a percentage of building limit or a higher flat dollar amount when a storm is named or crosses a trigger latitude. These deductibles can exceed \$100,000 on coastal commercial schedules and apply separately per location or per occurrence depending on form language. Verify whether the trigger uses National Weather Service naming, declarations from state insurance departments, or carrier-defined thresholds.

Wind-driven rain disputes turn on whether damage required a direct opening or leakage through an intact envelope. Ordinance-or-law, extended restoration, and debris removal sublimits shape net recovery—document pre-storm roof age and opening protection. For valuation fights after a large wind event, compare agreed value vs ACV and RCV.

Earthquake peril analysis

Physical characteristics and loss drivers

Earthquake loss is dominated by peak ground acceleration, duration, soil amplification, and building period. Soft-story wood frames, unreinforced masonry, non-ductile concrete frames, and poorly anchored parapets produce the familiar streetscape failures in Western states. Content damage, sprinkler leakage, fire following shaking, and lifeline outage (water, gas, power) add loss layers beyond structural drift.

Liquefaction and lateral spreading destroy foundations in saturated soils; hillside homes fail when retaining walls are under-designed. Aftershocks can re-damage weakened connections—plan for it in continuity work.

Building code and construction vulnerabilities

Code era is a proxy for ductility and connection detail. Pre-1970s concrete, brick veneer over flexible frames, and tuck-under parking soft stories are recurring retrofit targets. Bolt-down of cripple walls, shear panel installation, and chimney bracing are cost-effective residential mitigations. Critical facilities often require base isolation—but many mid-century commercial cores stay seismically deficient until a lender or tenant demands retrofit.

Coverage and underwriting issues

Standard property policies **exclude earthquake** unless bought back via endorsement or a standalone earthquake policy. Earthquake coverage uses specialized deductibles—commonly 5% to 25% of building value per location—and may sublimit masonry veneer, sprinkler leakage, and business income. Difference in conditions (DIC) policies can stack earthquake on top of a restricted primary form, but coordination of deductibles and follow-form language must be read carefully.

California Earthquake Authority (CEA) policies and private EQ markets differ on contents, loss of use, and code upgrade coverage. Mortgage holders may require EQ insurance in high-hazard zones even when not legally mandated. Brokers should map geotechnical reports and known fault proximity to realistic deductible exposure, not just premium. Gap analysis belongs in property insurance exclusions and endorsements.

Wildfire peril analysis

Physical characteristics and loss drivers

Wildfire destroys structures through radiant heat, direct flame contact, and ember intrusion—often embers hours ahead of the main front. Loss severity correlates with fuel continuity to the structure, vent screening, deck and fence attachments, roof valley debris, and adjacent outbuildings. Urban conflagration events show that dense WUI neighborhoods fail as systems: one ignited structure raises exposure for the entire block.

Smoke and ash trigger sizable contents and time-element claims even without total loss. Utility shutoffs and evacuations extend business interruption far from the fire perimeter.

Building code and mitigation vulnerabilities

Class A roofing, ember-resistant vents, tempered glazing, non-combustible siding in the zero-to-five-foot zone, and screened eaves reduce ignition probability. California and other WUI states incorporate Chapter 7A-style ember and exposure standards; local defensible-space ordinances require vegetation management and clearance around structures. Hardening without clearing fuels yields partial benefit; insurers ask for photos, third-party wildfire scores, and mitigation certificates at bind and renewal.

Coverage and underwriting issues

Admitted carriers may non-renew or cap limits in high fire severity zones, pushing owners to surplus lines or the California FAIR Plan—a basic fire insurance pool of last resort. In April 2026 regulators approved an average **29.1% dwelling fire rate increase** for the FAIR Plan effective October 15, 2026; individual changes depend heavily on the wildfire portion of premium and hardening discounts. FAIR Plan growth slowed to about 1.9% in Q2 2026 as some standard-market capacity returned to distressed areas, per California Department of Insurance alerts—but FAIR Plan plus a separate DIC policy remains a common stack for full homeowners-equivalent coverage.

Wildfire deductibles may be flat dollar or percentage-based. Underwriters examine roof age and material, tree overhang, access for fire apparatus, and prior claims within a mile radius. Community-level mitigation (fuel breaks, undergrounding) is beginning to influence renewal outcomes but property-level defensible space is still the lever owners control directly.

Cross-peril program design for owners and risk managers

Align limits with realistic rebuild cost inflation in high-demand post-event labor markets. Separate wind, flood, earthquake, and fire deductibles can stack in a single calendar year if storms and quakes coincide with fire season—stress-test combined out-of-pocket exposure.

Keep an asset register with roof install date, retrofit documents, elevation certificates for flood, and geotech references for EQ sites. Catastrophe response playbooks should name peril-specific adjusters and engineers; wind and EQ disputes often turn on code upgrade and ordinance coverage. For claim execution, use claims management: the complete professional guide.

Reinsurance and ILS capacity for U.S. wind, quake, and fire risk remained robust in the first half of 2026—Artemis tracked roughly **\$17.98 billion** of catastrophe bond issuance across **83 transactions** in H1 2026, with outstanding cat bond market size near \$65.6 billion at mid-year. That capital supports carrier limit but does not replace peril-specific sublimits or exclusions on your dec page. See parametric insurance and catastrophe bonds for how index-based tools complement traditional indemnity covers.

Accumulation control—too many values in one surge footprint, fault segment, or wildfire corridor—drives renewal pricing and capacity at the portfolio level. Risk managers placing multi-state schedules should read catastrophe portfolio management alongside this peril breakdown.

Frequently Asked Questions

Does standard property insurance cover hurricane wind and flood together?

Most property forms cover hurricane wind unless limited by endorsement, but flood—including storm surge—is excluded and must be bought through the NFIP or private flood insurers. Under Risk Rating 2.0, NFIP premiums reflect individual flood characteristics, not zone averages. Named-storm deductibles may apply to wind separately from any flood policy deductible.

Why is earthquake damage excluded on my property policy?

Earthquake is catastrophic and geographically concentrated, so standard forms exclude it—buy an endorsement or standalone earthquake policy. Those policies use high percentage deductibles and specialized limits. Without them, shaking, fire following earthquake, and seismic sprinkler leakage may be uninsured—read exclusions

alongside any DIC policy.

What triggers a hurricane or named-storm deductible?

Triggers vary by carrier and state form: a named tropical cyclone, a NOAA or state insurance department declaration, or wind speed at the location. The deductible is often a percentage of building limit—5% to 10% on residential coastal policies, higher on commercial schedules. Confirm whether it applies once per season or per event, and whether locations aggregate.

How does defensible space affect wildfire insurance?

Insurers reward reduced fuel in the zero-to-five-foot zone, cleared roofs and gutters, ember-resistant vents, and removal of combustible attachments like fences to decks. Many carriers request photos or inspections at renewal. Mitigation lowers ignition risk but does not guarantee renewal if community wildfire scores or carrier ZIP-code aggregates are exceeded.

Is the California FAIR Plan the same as full homeowners insurance?

No. The FAIR Plan provides basic fire coverage when admitted markets decline the risk—not a full HO-3. Owners typically pair it with a DIC policy for liability, water, theft, and loss of use. Rates track wildfire loss experience—2026 increases hit the wildfire portion of premium hardest. Shop admitted and surplus lines before accepting last-resort placement.

Can one property need separate actions for all three perils?

Yes. A coastal California or Gulf property may need wind coverage with a named-storm deductible, NFIP or private flood for surge, earthquake buy-back, and wildfire hardening. Each peril needs different retrofits and policy parts. A single blanket limit without peril-specific deductibles and exclusions understates true out-of-pocket risk.

Catastrophe Modeling: How RMS, AIR, and Verisk Quantify Catastrophe Risk

Source: <https://riskcoveragehub.com/catastrophe-modeling-rms-air-verisk-quantify-risk/> Published: March 17, 2026

Updated October 6, 2026.

Catastrophe models from RMS, AIR, and Verisk translate physical peril science into probabilistic loss estimates insurers use for pricing, reinsurance, and accumulation control. Each vendor builds stochastic event catalogs, damage functions, and financial engines that roll up to metrics such as average annual loss (AAL), probable maximum loss (PML) at chosen return periods, and occurrence and aggregate exceedance probability curves. The outputs are long-horizon views of tail risk that must be validated, calibrated to experience, and read alongside underwriting judgment.

Where RMS, AIR, and Verisk sit in the market

For decades, property cat risk in North America and many global markets has been shaped by two large independent model families: Risk Management Solutions (RMS) and AIR Worldwide (AIR). Both were acquired by Verisk and now sit under a single analytics umbrella while retaining distinct model lineages, data pipelines, and client bases. Many carriers run one vendor as a primary view, the other as a challenger, or blend

outputs in internal aggregation systems.

Practical differences show up in event catalogs, peril coverage, vulnerability assumptions, and financial layering—not in the abstract goal. Each vendor documents peril modules (hurricane, earthquake, severe convective storm, wildfire, flood where offered, and other secondary perils), update cadence, and regional footprints. Underwriters and catastrophe analysts choose a vendor—or a multi-vendor stack—based on portfolio mix, reinsurance program design, and how well model losses reconcile with their own claims history.

Risk managers who do not run models in-house still encounter these tools through insurer submissions, broker analytics, and treaty pricing. Understanding what the vendor is actually quantifying prevents mistaking a modeled PML for a policy limit or a single scenario for a full tail distribution. For broader context, see the hub's complete guide to catastrophe modeling and how modeled climate trends feed pricing.

How a catastrophe model is built

A cat model is three linked engines plus a roll-up layer. Vendors publish methodology papers; implementation detail varies by peril and region, but the architecture is consistent enough that risk teams can audit each stage.

Hazard and stochastic event generation

The hazard module defines how events are drawn from history, physics, and statistical extrapolation. For hurricanes, that typically includes track, intensity, and landfall logic tied to basin climatology. Earthquake modules use fault and seismicity models. Severe convective storm and wildfire modules use different generators—hail and tornado footprints versus fire spread and ignition patterns tied to fuels, weather, and topography.

The output of this stage is a large stochastic catalog: thousands to hundreds of thousands of events, each with frequency and spatial footprint. Event rates and severities are tuned so that modeled distributions align with historical catalogs and, where vendors offer them, forward-looking or climate-conditioned views that adjust selected peril parameters.

Vulnerability and exposure

Vulnerability functions (damage curves) map hazard intensity at a location to a damage ratio or damage state for a given construction and occupancy class. Exposure inputs describe values at risk: replacement cost, business interruption, contents, deductibles, limits, and secondary modifiers such as year built, roof type, or mitigation features when captured.

Small shifts in vulnerability assumptions can move tail losses materially for high-gradient perils (hurricane wind and storm surge near coasts, earthquake shaking in soft soil, wildfire in the wildland–urban interface). That is why model change management is a standing agenda item at renewals. Property owners see the other side of the same exposure data in policy structure—how Coverage A, B, C, and D apply.

Financial module and metrics engine

The financial module applies policy terms to ground-up damage: deductibles, limits, coinsurance, sublimits, waiting periods on time element coverage, and reinsurance in gross versus net runs. From simulated event losses, the engine builds occurrence and aggregate distributions and summary metrics.

Occurrence loss groups all locations hit by one event before policy terms. Aggregate loss sums all events in a period (often a calendar or treaty year). That distinction drives whether a program responds on a per-event excess-of-loss tower or an aggregate stop-loss layer—topics covered in depth in the hub’s reinsurance treaty guide and quota share versus excess-of-loss structures.

Outputs risk teams actually use

Average annual loss (AAL)

AAL is the mean loss per year across the simulated catalog, net of terms and optionally net of reinsurance. It is a central tendency metric useful for pricing expected loss load and comparing portfolios over time. AAL alone says little about tail risk; two portfolios with similar AAL can have very different PMLs if one concentrates coastal wind or wildfire exposure.

PML and return periods

Probable maximum loss (PML) in industry practice usually means a high quantile of the modeled loss distribution at a stated return period—often 100-year, 250-year, or 500-year occurrence, depending on peril and internal standard. The label “PML” is not interchangeable across organizations unless return period, occurrence versus aggregate basis, gross versus net, and peril scope are defined.

Occurrence and aggregate exceedance probability curves

An occurrence exceedance probability (OEP) curve plots the probability that a single event exceeds a loss threshold. An aggregate exceedance probability (AEP) curve plots the probability that total annual losses exceed a threshold. Reinsurance buyers use OEP views to size event caps; capital and aggregate treaty discussions lean on AEP. Tail Value at Risk (TVaR) captures the average severity of outcomes worse than a chosen PML point—not just the threshold loss.

Gross versus net and model blending

Insurers run gross models to design reinsurance, then net models to reflect ceded recoveries under treaty terms modeled with varying fidelity. Multi-vendor shops may take ensembled or bounded views when RMS and AIR diverge on the same book. The discipline is to document which view anchors pricing versus which view stress-tests the program.

Validation, calibration, and model change

Models are hypotheses backed by data and expert judgment. Validation compares modeled loss distributions and event rates to independent benchmarks: historical industry losses, company experience where credible, and peril-specific studies. Calibration adjusts components—event rates, vulnerability anchors, or financial flags—so that modeled AAL and selected return periods sit inside agreed tolerance bands.

When vendors release major updates, clients rerun baseline portfolios and trace drivers: new hazard science, exposure schema changes, vulnerability retuning, or financial engine fixes. Underwriting and actuarial teams align on whether to phase changes into pricing or absorb them in trend factors. Secondary perils that once sat in “other natural catastrophe” buckets now get explicit modules; that shift alone can reorder rank by peril for inland books.

Honest use of cat output treats the model as one input. It does not replace field underwriting, inspection data, or knowledge of local building code enforcement. It also does not replace reading the policy—modeled limits must match bound coverage, including exclusions that standard property forms leave open.

2026 model themes: climate-conditioned views and secondary perils

Vendor release notes in recent cycles emphasize two parallel tracks: forward-looking hazard views that reflect selected climate science pathways for certain perils and regions, and expanded treatment of perils that drive growing attritional and tail loss outside traditional hurricane and earthquake zones.

Climate-conditioned hurricane and flood-related parameters, where offered, adjust aspects of severity or frequency relative to baseline catalogs. Users choose whether to run baseline, alternate, or blended views; the important operational point is to label results clearly in reinsurance submissions and internal capital discussions so year-over-year comparisons are not apples-to-oranges.

Secondary and tertiary perils—severe convective storm (hail, tornado, straight-line wind), wildfire, and inland flood where modeled—now receive dedicated event generators and vulnerability suites on major platforms. That matters for mid-continent carriers and any account with wildland–urban interface exposure. It also shapes property pricing when modeled tail risk rises even if recent loss history is quiet—see the hub’s overview of hard versus soft market cycles and coverage pricing.

Earthquake, hurricane, and wildfire modules continue to receive regional refinements (landfall logic, fire spread, liquefaction and site amplification where applicable). Risk managers should read vendor release documentation for their perils and territories rather than relying on headline press summaries.

How insurers and risk managers apply the outputs

Property pricing and portfolio steering

Commercial and personal lines insurers use modeled AAL and tail metrics to set rate adequacy targets by territory and construction class, cap aggregate deployment in high-return-period zones, and justify deductibles or sublimits. Large insureds with scheduled locations may see cat analytics reflected in capacity placement and minimum attachment points.

Reinsurance structure and renewal

Reinsurance brokers and ceding insurers use occurrence EP curves to layer event excess-of-loss programs and AEP views to size aggregate covers. Modeled net PML after reinsurance supports discussions of retention, co-participation, and reinstatement economics. When modeled losses shift at a fixed return period, attachment points and rates on excess layers move even if the prior year’s experience was benign.

Accumulation management

Cat models power aggregation by event: sum insured within a hurricane footprint, fault rupture scenario, or wildfire burn probability swath. Risk teams set caps by zone, monitor roll-ups against tolerances, and stop binding when modeled aggregate PML or TVaR breaches guidelines. Corporate risk managers use the same logic indirectly when insurers aggregate locations on declarations and require updated schedules after acquisitions or new construction.

Corporate and captive perspectives

Organizations retaining risk through deductibles, captives, or parametric triggers still live in a market priced with vendor cat analytics. Modeled tail loss informs retention sizing, collateral discussions, and whether to buy incremental limit above primary. The same renewal discipline applies: know what the model assumes about values and locations.

Frequently asked questions

What is the practical difference between RMS and AIR models?

Both vendors produce stochastic event catalogs, vulnerability functions, and financial roll-ups, but they differ in peril modules, historical data blends, and damage assumptions. Many insurers treat one as primary and the other as a benchmark; large divergences at a chosen return period usually trace to hazard footprints or vulnerability curves, not to a single “correct” number.

How should I interpret AAL versus PML?

AAL is the average modeled loss per year and supports expected-loss pricing. PML is a high quantile at a stated return period on an occurrence or aggregate basis and describes tail risk. Use them together: AAL for central tendency, PML and EP curves for reinsurance and capacity limits.

When do I use an occurrence EP curve instead of an aggregate EP curve?

Use the occurrence EP curve when a single event drives recovery—typical for per-event excess-of-loss reinsurance. Use the aggregate EP curve when multiple events in one treaty year stack toward an aggregate limit or stop-loss. Mixing bases without relabeling is a common source of mis-sized towers.

What does model validation and calibration actually check?

Validation compares modeled loss distributions and event rates to credible external references and company experience where data allows. Calibration adjusts model components so agreed metrics—often AAL and selected return periods—fall within tolerance. It does not guarantee the next event will match the catalog; it keeps the tool honest against history and stated assumptions.

What changed in recent vendor updates around climate and secondary perils?

Release cycles have added or expanded climate-conditioned hazard views for selected perils and regions, plus dedicated severe convective storm and wildfire modules with their own event generators. Users should note which view they ran—baseline versus forward-looking—and expect secondary peril contributions to rank higher on inland and wildland–urban interface books.

How do cat models feed reinsurance buying decisions?

Modeled occurrence losses size event excess layers; aggregate losses size stop-loss and aggregate covers. Net runs after treaty terms show remaining tail retention. Brokers and cedents anchor attachments and limits to EP points and TVaR where internal policy requires, then negotiate rate and capacity against that modeled net profile.