

OPEN CLIMATE RISK FRAMEWORK · OCRF V1.0 · COMPENDIUM

The Climate Risk Commons

Why we are opening bank-grade climate-risk infrastructure to students, researchers, and the institutions the market forgot.

Climate financial risk has no Linux. We built it. Climate Risk Commons puts the same enterprise-grade physical climate-risk platform that the world's largest banks use into the hands of any researcher, free for non-commercial work: not only standardized data, but the platform to upload open datasets, run experiments, build new models, and design new damage functions.

A not-for-profit, founded and funded by RiskThinking.AI · riskthinking.org

CONTENTS

Institutions must report numbers they cannot defend

The mandate is now both real and prudential. Measuring physical climate risk has shifted from a voluntary exercise to a supervised obligation. IFRS S1 and S2, the CSRD, the EU Taxonomy, TCFD and TNFD have made climate and nature risk a required part of corporate and financial reporting, and prudential supervisors from the EBA to BaFin and FINMA now expect institutions to identify, measure, manage, and monitor physical climate-risk drivers.

The harder test is ownership, not reporting. Auditors and supervisors increasingly expect an institution to understand and own the data and models behind a tool, rather than accept a vendor's black box. A closed, proprietary model cannot meet that bar by design.

The market is opaque, and the numbers do not agree. Independent reviews of the field, including UNEP FI's Climate Risk Landscape work, keep identifying the same compounding problems: provider methodologies that differ substantially and remain largely opaque, so results cannot be benchmarked.

A single vendor manufactures false precision. The same asset, scored by two providers, can yield materially different risk scores. Relying on any one closed vendor produces a false sense of precision, exactly when supervisors are asking institutions to understand model uncertainty rather than look past it.

An open standard is the structural cure. Inspectable methodology lets an institution own its models. A shared schema makes results comparable across teams and firms. A stochastic, full-distribution engine makes uncertainty legible. Open, no-cost access reaches the smaller institutions, emerging markets, and researchers the commercial market has priced out.

The timing is decisive. Climate science has just moved to CMIP7, replacing CMIP6 whose data stopped at 2014. Climate Risk Commons is among the first production systems to ship on CMIP7, so members start on the current standard rather than last decade's.

WHAT IT IS

An open standard, anchored by two pieces

Climate Risk Commons is an open technical standard for physical climate-risk analysis. The Open Climate Risk Framework (OCRf) defines the shared contracts, asset schema, reference pipeline, APIs, SDKs, and tooling that let hazards, assets, and impact functions interoperate. The reference pipeline, Ecofusion, runs the entire specification end to end on a laptop or a cluster using open reference data. Every component is licensed under Apache 2.0, with terms you can hand to your legal office. Through CDTexpress, members reach the full Climate Digital Twin engine behind it: 50+ climate hazards, 241 billion geospatially aligned locations, 193 countries, all IPCC emissions scenarios, and 15 time horizons to 2100.

This is Climate 2.0: stochastic, multi-hazard, coherent across hazards and countries, and built for the extreme tail rather than the average case.

Eight components of membership


- **OCRf spec and asset schema.** How assets and exposure are encoded, so analyses share a common vocabulary.
- **Ecofusion reference pipeline.** An end-to-end DataFusion engine that runs on a laptop or a cluster.
- **Impact-function library.** Validated functions for common asset-hazard pairs.
- **SDK and API clients.** Permissive Python and R clients as the on-ramp to the standard.
- **Sample notebooks.** Portfolio analysis, sensitivity testing, and regulatory submissions.
- **Conformance test suite.** What an implementation must pass to claim conformance, so outputs remain comparable.
- **Documentation.** Workflows, uncertainty explainers, and regulatory mapping.
- **Multi-hazard stress testing.** An asset's exposure to several correlated hazards at once.

What you can build

- **See your real exposure.** Load asset or loan-book addresses; the twin geolocates them and returns flood, heat, and wildfire exposure for the assets actually held, not a headline figure.
- **Stand up a CMIP7 stress test.** Run a stochastic, multi-pathway test across thousands of scenarios, in a structure a supervisor can read and an auditor can follow.
- **Bring your own data.** Load your own datasets, models, and damage functions and still produce a result that compares cleanly with everyone else's.
- **Screen a city or development portfolio.** Development banks, ministries, and city networks screen assets for priority physical risks and share the method with anyone who needs to use it.

HOW IT WORKS

Three open contracts, one comparable result



Anyone can run the standard end to end on open reference data. Swap in production-grade inputs when a decision depends on it.

WHAT YOU DO NOT GET

An open standard earns trust by being explicit about its boundaries. Here are ours.

Membership provides the engine, not our commercial data. Members receive the full machinery for calculating physical climate risk, together with an environment for storing open-source datasets that can be downloaded directly to power it. Membership does not include RiskThinking.AI's proprietary high-fidelity data. That data is expensive to build and maintain, and it is reserved for paying customers. The commercial revenue it earns is what sustains the open platform.

The engine is the same one the largest institutions run. Nothing in the Commons is a demonstration version. Members work with the same engine that powers the world's largest financial institutions, applied to open reference data or to whatever data, models, and damage functions they bring themselves. That makes the Commons an excellent base for developing and testing new, scientifically grounded ideas.

The advantage is comparability, not a dataset. The primary advantage researchers gain from the Commons is the ability to share and compare results built on well-developed standards. Two teams using the same schema, the same conformance suite, and the same engine can replicate one another's work and compare findings directly. No single dataset, however good, delivers that. Only a standard does.

The Commons is a not-for-profit, founded and funded by RiskThinking.AI, using the same standards as the commercial platform. That shared foundation keeps open research and commercial practice interoperable rather than letting them drift apart.

Governance built to outlast us

Extend the frontier the market has skipped. Commercial vendors cluster around large, listed, global portfolios. The underserved long tail, including agriculture and supply chains, real estate and unlisted assets, single river basins, and specific hazards and geographies, is exactly where researchers work. Building that science on the open standard adds the coverage the market has left out and folds it into a shared, comparable foundation rather than another silo. Researchers are the people who pressure-test a standard, validate it, and extend it into the underserved corners, giving it a legitimacy no single vendor's marketing can buy.

Shape something built to outlast any vendor. RiskThinking is the first maintainer of the standard, not its owner. As adoption grows, a neutral foundation follows, so what members build on remains in place in five years. Members retain full freedom to bring their own data, models, and damage functions, while holding their work to a standard the whole community shares.

Roadmap

- **Groundwork.** Resolve the patent-rights path, select the licence, prepare the repository, finalize spec v1, and onboard pilot users.
- **Ecosystem ignition.** Public release of Ecofusion, OCRF, and the schema; the open tier goes live; design partners, including universities and innovation labs, join.
- **Standard expansion.** Community contributions, academic partners, deeper regulatory engagement, and the transition to a neutral foundation.
- **Scale and sustain.** Premium data, support, and advanced capabilities sustain the open core, keeping the standard continually improving for everyone building on it.

Expect rough edges. Specs, APIs, and data schemas may change during the preview, sometimes with little notice. Availability is a plan, not a promise; the exact order and timing of drops may shift.

INSTITUTIONS ALREADY ON BOARD

The Fields Institute · University of Toronto (Architecture) · University of Toronto (Mathematical Finance) · Tel Aviv University · Durham University (Finance and Management Science) · The Cyprus Academy of Sciences, Letters and Arts · Pacific Institute for Climate Solutions, University of Victoria · Technical University of Munich (CAMBIR) · University of Western Ontario (All Science Faculties) · London School of Economics (doctoral research) · University of Zurich · Queen's University, Smith School of Business, Institute for Sustainable Finance · Central Bank of Brazil · Universidade Federal do Rio Grande do Norte (UFRN) · Stockholm Environment Institute (SEI Latin America) · Twintree.Org · University of British Columbia, Sauder School of Business · University of Duisburg-Essen · Jewish Climate Trust

Questions we hear from partners

01 What is the Climate Risk Commons?

An open initiative to give the climate physical-risk field a shared technical foundation. Today, climate-risk analysis is fragmented across data vendors, asset taxonomies, hazard models, impact functions, and reporting workflows: outputs are hard to compare, assumptions are hard to audit, and providers are hard to integrate. The Commons addresses this by publishing an open standard, a reference pipeline, open data, and tooling that anyone can inspect, run, and build on.

It is a not-for-profit, founded and funded by RiskThinking.AI, with governance moving to a neutral foundation so the standard outlives any single company's commercial decisions.

02 What's the difference between the Commons and OCRF?

The Climate Risk Commons is the initiative and the ecosystem: the community of partners, the open assets, and the governance that stewards them. The Open Climate Risk Framework (OCRF) is the core technical specification at its heart, defining common contracts for the three inputs to physical-risk analysis: climate hazard data, physical assets, and impact functions. A simple way to say it: the Commons is the community and the home; OCRF is the standard the community builds around.

03 What's included in the Commons?

The open assets, free to inspect, run, and build on: the OCRF specification with input contracts and the output model; Ecofusion, an open reference pipeline implementing OCRF end to end; open reference data to run and validate without a commercial relationship; a starter library of open impact (damage) functions; output APIs and SDKs plus tooling and documentation; and open governance with spec versioning, a published contribution process, and a public roadmap. Production-grade results at scale typically draw on commercial data and functions, from RiskThinking.AI or other conforming providers.

04 Why use OCRF?

Comparability: when hazards, assets, and impact functions conform to the same contracts, results from different vendors, models, and internal teams can be placed side by side and meaningfully compared. Auditability: every assumption is inspectable, and tracing an output through an open pipeline beats black-box scores. No lock-in: open contracts let you swap providers, mix data sources, or bring your own functions without rebuilding your workflow. Lower cost to start: the spec, Ecofusion, reference data, and starter functions are free to run.

The pitch in one line: OCRF turns climate physical-risk analysis from incompatible vendor black boxes into an open architecture you can inspect, compare, and own.

05 What should the ecosystem build?

Conforming data products published against the OCRF input contracts, including from providers who compete with RiskThinking.AI; sector-, peril-, and region-specific impact functions, open or commercial; integrations into risk platforms, portfolio systems, GIS tools, reporting workflows, and cloud data platforms; alternative implementations, validation and conformance tooling, and benchmarking suites; research that tests and

extends the methods encoded in the spec; and applications for underwriting, lending, asset management, supply chain, infrastructure planning, and regulatory reporting.

06 What does building in the ecosystem look like?

Orient: read the OCRF spec and the three input contracts, clone Ecofusion, and run it end to end on the reference data. Pick your surface: data providers publish against an input contract, modelers contribute or sell impact functions, tool builders consume outputs through the APIs, implementers build specialized pipelines. Build against the contracts, using the reference pipeline and open data as your test harness. Validate conformance, which makes your product plug-compatible with everything else in the ecosystem. Contribute back where it makes sense. Ship and distribute: every OCRF adopter is a prospective user of every conforming component.

07 How do academics, researchers, and developers benefit?

Researchers and academics gain reproducibility by default, since methods published against OCRF re-run by anyone on open data in an open pipeline; a path from paper to practice as contributed methods get adopted by practitioners; shared infrastructure that removes months of setup; and influence on how v1 encodes hazards, exposure, and vulnerability.

Developers and implementers gain a stable foundation of open contracts and versioned specs; distribution through the standard, since every adopter is a prospective user; a zero-friction start, with pipeline, sample data, SDKs, and docs free to run; and community visibility with a direct line to the team and a voice in the roadmap.

08 Do members get RiskThinking.AI's commercial data?

No. RiskThinking.AI's high-fidelity data is reserved for paying customers. Members get the machinery for calculating climate risk and an environment stocked with open-source datasets they can easily download to power it. What differs is the data flowing through it, not the machinery doing the work.

09 Is the engine the same one commercial clients use?

Yes. The Commons runs on the same engine that powers the world's largest institutions. Nothing in the Commons is a demonstration version.

10 What data can members use?

The open reference datasets provided within the Commons environment, any open dataset members upload, and any data, models, or damage functions of their own. Results remain conformant with the standard regardless of whose data powers them.

11 Are results produced on open data scientifically credible?

Yes. The engine, schema, and conformance suite are identical to the commercial stack. Precision depends on the resolution of the data used, and members should report that, exactly as good science already requires.

12 Why open the machinery but not the data?

Because the field's structural problem is the absence of a shared standard, not a shortage of datasets. Opening the machinery fixes comparability for everyone. The high-fidelity data is costly to build and maintain, and the commercial revenue it earns is what sustains the open platform.

13 Who owns and runs the Commons? Can I trust its neutrality?

Climate Risk Commons is a not-for-profit, founded and funded by RiskThinking.AI, and it uses the same standards as the commercial platform. Governance is lightweight but real: clear spec versioning, a published contribution process, reference implementations, and a public roadmap. RiskThinking is the standard's first maintainer, not its owner; as adoption grows, a neutral foundation follows. Competing providers are explicitly welcome to conform and contribute.

14 What if a member later needs high-fidelity data?

That is a commercial conversation, and we are happy to have it. Write to academic@riskthinking.ai and we will route it to the right team.

15 Can membership be used for commercial work?

No. Membership is free for non-commercial research only. Commercial use of the platform or of RiskThinking.AI's data requires a commercial licence.

16 How do I get involved?

Join the early access preview: send your GitHub handle and contact email and confirm the early access terms. We'll onboard you to the repositories, the community channel, and the spec. From there, run the reference pipeline on the sample data and tell us what you find: early feedback directly shapes the public v1 release.

HOW TO JOIN

Membership is free for non-commercial research

- **Individual researchers** request access by writing to academic@riskthinking.ai, with proof of institution and a sentence or two on intended use.
- **Faculties and universities** can enrol all their bona fide students, researchers, and faculty at once. A short letter of intent from the institution is enough to formalize it.

The invitation. Closed competitors are racing to build a wall. We are building a community. It is ready for you to read, run, and build on today. Join the community that evolves this project. Contribute, and gain from the contributions of others.

REQUEST ACCESS academic@riskthinking.ai

OPEN SOURCE COMMUNITY opensource@riskthinking.ai

WEB riskthinking.org