BIS Innovation Hub



Project Danu

Monitoring emerging risks by leveraging the concept of a digital twin

November 2026

Banque de France Eurosystème
De Nederlandsche Bank Eurosysteem
Hong Kong Monetary Authority 香港金融管理局
MAS Monetary Authority of Singapore

© Bank for International Settlements 2026.

All rights reserved.

Limited extracts may be reproduced or translated provided the source is stated.

ISBN 978-92-9279-004-2

Executive summary

Climate-related physical risks are an increasingly important concern for central banks and supervisory authorities due to their potential impact on financial stability. Extreme weather events such as cyclones, droughts, floods and wildfires can disrupt households, businesses, infrastructure and financial institutions, with implications for both economic activity and financial stability.

Assessing these risks remains challenging. Central banks and supervisory authorities often face limitations in data availability, analytical tools and real-time monitoring capabilities. A particular challenge is that relevant climate-related information is often distributed across multiple external data sets that differ in format, coverage and quality. This makes it difficult to develop a consistent and timely representation of how physical hazards affect assets, economic activity and the financial system.

Project Danu explored how the concept of a digital twin can be made practicable for climate-related risk analysis. A digital twin is a virtual representation of a real-world system. Building on a proof of concept developed by the Bank of France, De Nederlandsche Bank and the Hong Kong Monetary Authority, the key question was how the approach can be made scalable, adaptable and reusable by central banks and supervisory authorities for a broad range of risk analysis use cases.

The project explored key design choices and verified these choices in a prototype. It identified that the combination of a structured data architecture and modular analytical components can support adaptability and scalability, while the use of open source technologies facilitates integration across heterogeneous technical environments.

The resulting prototype integrates environmental, geospatial, asset and financial information within a common analytical framework, linking climate-related hazards to financial impacts through the hazard-exposure-vulnerability-finance methodology commonly used in catastrophe modelling. The framework can be extended with additional data sources, hazards and analytical capabilities to support a broader range of climate risk use cases.

The results provide a pathway for central banks and supervisory authorities seeking to generate timely indicators of the potential impacts of extreme weather events and support the assessment of risks to financial institutions and financial stability. The project establishes a foundation for future development towards the incorporation of predictive and simulation capabilities as well as additional damage and financial models to support more forward-looking climate risk analysis.

Contents

Executive summary 3

Climate-related physical risks and the financial system 5

Project overview 7
Project Danu 7
Previous work 8

Framework implementation and scalability 9
Framework design 9
User roles 10

Data and methodology 11

Case study and results 14
Insights from the case study 15

Limitations and potential next steps 16
Limitations 16
Potential next steps 17

Summary 17

Project partners and acknowledgements 18

Annex 20

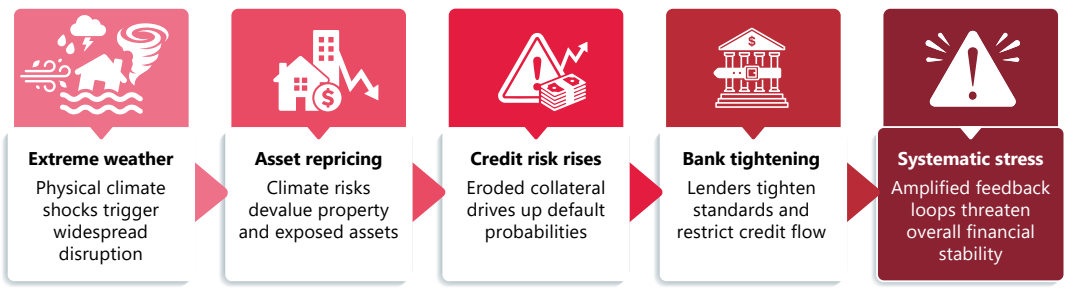
Climate-related physical risks and the financial system

Emerging risks are evolving threats that may have significant implications for the financial system. Their transmission channels, magnitude or likelihood remain only partially understood. Project Danu focused specifically on climate-related risks, particularly physical hazards, which are considered emerging due to their evolving nature, uncertain impact and limited precedent.

There is growing urgency to monitor climate-related risks within the financial system. Physical climate hazards such as floods, wildfires, heatwaves and rising sea levels are increasingly recognised as material risks with significant macroeconomic and microeconomic implications. As part of their financial stability mandate, central banks and supervisory authorities must develop a deeper understanding of how such risks spread through the financial system.

As climate-related risks increasingly materialise, physical hazards can affect the financial system through several channels. Direct damage to homes, businesses and infrastructure can reduce asset and collateral values, leading to the repricing of climate-exposed assets and raising credit risk for lenders. Extreme weather events can also generate significant insurance and underwriting losses, affect the availability and affordability of insurance coverage, and transmit risks through reinsurance markets. In addition, disruptions to economic activity, critical infrastructure and financial services can create liquidity and operational challenges for households, firms and financial institutions. Where these impacts are sufficiently large or widespread, they may weaken economic activity, affect public finances and generate broader financial stability risks through interconnected exposures and feedback loops between the financial sector and the real economy (Figure 1).

Illustrative impact of physical climate risks to financial stability — Figure 1



The examples in Box A show how extreme weather events and natural disasters can generate significant impacts that extend beyond the physical economy and into the financial system. Understanding and monitoring these impacts are therefore becoming increasingly important for central banks and supervisory authorities in fulfilling their financial stability mandates. However, institutions often face challenges related to data availability, modelling capabilities and the generation of timely and actionable insights – gaps the Network for Greening the Financial System (2026)[1] identifies as a key priority for future work.

A particular challenge is that relevant climate-related information is often distributed across multiple external data sets that differ in format, coverage and quality. While the availability of environmental and geospatial data continues to improve, integrating these data sets and translating them into timely and actionable insights remains resource-intensive. In addition, many existing approaches are primarily designed for periodic assessments and can be less suited to analysing rapidly evolving events as they unfold. These challenges highlight the need for approaches that can integrate diverse data sources and support more dynamic forms of climate risk analysis.

Box A: Recent examples of climate risk transmission

Recent extreme weather events illustrate the channels through which such events affect financial systems. Examples include the 2024 floods in Valencia, Spain, and the 2025 wildfires in Los Angeles, United States.

The September 2024 floods in eastern Spain illustrate several risk channels simultaneously. Credit risk materialised as banks estimated that €20 billion in loan exposure in the affected regions came under strain when borrowers temporarily lost income or productive assets. Liquidity risk emerged as payment infrastructure and cash access were disrupted, while underwriting risk was borne by insurers facing large claims. In response, banks implemented payment holidays and loan restructuring, and the Bank of Spain monitored exposures to ensure continued access to essential financial services. Drawing on 31 case studies compiled directly from central banks across 28 economies, the Network for Greening the Financial System (NGFS) (2026) shows that this same cluster of credit, collateral and liquidity risks recurs across very different jurisdictions. The NGFS attributes Spain's contained outcome to the depth of its public-private insurance framework, which absorbed a substantial share of losses and dampened the transmission from physical damage to bank balance sheets. While the banking system proved resilient, the event demonstrated how climate shocks can rapidly translate into financial shocks.

The 2025 California wildfires point to a related but distinct channel that the NGFS identifies as central to how climate risk now spreads through the financial system: market and underwriting risk repricing transmitted through reinsurance and collateral markets rather than through direct losses to financial institutions. Although insurers remained solvent, the event accelerated the repricing of wildfire risk, pushing up

1. Network for Greening the Financial System (NGFS), NGFS Note on the economic and financial impacts of extreme weather events, May 2026, https://www.ngfs.net/en/press-release/ngfs-publishes-report-macroeconomic-and-financial-impact-extreme-weather-events.

reinsurance costs, prompting insurers to withdraw from high-risk areas and raising concerns over mortgage collateral values as properties become uninsurable, harder to insure or lose value outright. The NGFS note explains this as a spillover mechanism: because national insurers transfer a significant share of disaster-related losses to international reinsurance markets, a localised hazard can generate financial pressure well beyond the affected region. Where insurance retreats rather than adapts, the residual risk does not disappear but shifts, either onto exposed households and firms directly or onto public balance sheets. This is the mechanism visible in California: as insurance becomes scarcer or costlier, banks' reliance on insured collateral means mortgage default probabilities rise and expected losses grow even without any physical damage to the institutions themselves, generating financial stability concerns purely through the repricing of risk.

Project overview

Project Danu

Project Danu explored how the digital twin[2] concept can be applied by central banks and supervisory authorities in monitoring and analysing climate related physical risks. The project developed a prototype framework that combines environmental, geospatial, asset and financial information within a common analytical environment to create a digital representation of how climate-related hazards affect physical assets and translate into financial impacts.

The project applied the key principles of the digital twin concept, including the integration of heterogeneous data sources, geospatial representation of physical and financial exposures, near real-time monitoring of evolving hazards and the incorporation of hazard scenario information within a common framework. Through this approach, the project explored how digital-twin-inspired architectures can support climate risk assessment in a manner that is practical, scalable and adaptable across jurisdictions. The resulting prototype establishes a foundation upon which more advanced simulation, forecasting and analytical capabilities can be developed over time.

A key outcome was the identification of design requirements to support practical implementation across institutions, including interoperability, scalability, adaptability to different institutional environments and the ability to support users with varying levels of technical expertise.

The project focused on transmission channels linked to asset damage, collateral values and economic losses. It investigated how physical hazards can be translated into indicators of economic and financial impact, while other channels, including insurance, liquidity and sovereign risk, remain outside scope. To support this analysis, the prototype integrated hazard, exposure, vulnerability and financial risk methodologies drawn from the climate risk and catastrophe-modelling literature.

2. Digital twin technology refers to a digital replica or virtual representation of a physical object, system or process, used to monitor, analyse and simulate the real-world counterpart in real time.

Rather than developing new hazard models, the project explored how these methodologies can be combined within a common framework to generate insights for central banks and supervisory authorities. The initial scope covered cyclones, droughts, floods and wildfires.

The project was developed in close collaboration with a user reference group of central banks in five continents. Their input helped shape the prototype and project scope, ensuring the project's design and outputs reflected the needs of central bank and supervisory authorities across jurisdictions.

Previous work

Project Danu builds on a series of previous initiatives of the Bank for International Settlements Innovation Hub (BISIH) that have explored the intersection of climate risk, data and financial systems.

Most recently, Project Gaia[3] developed a pipeline to extract climate-related performance indicators from publicly available corporate reports. By leveraging generative artificial intelligence, the project demonstrated how to reduce the manual effort required for climate data collection and analysis. In tandem, Project Symbiosis[4] used artificial intelligence to estimate the carbon footprint in supply chains and frontier models to match upstream suppliers with transition finance.

Before this, Project Ellipse[5] introduced a prototype data analytics platform aimed at enhancing the global regulatory community's ability to operate in a more proactive, data-driven manner. Building on this, Project Viridis[6] explored the integration of climate risk data into such an analytics platform. In addition, in collaboration with the NGFS, the BISIH developed Project NGFS Data Directory 2.0,[7] addressing data gaps through a comprehensive mapping of available climate-related data sets.

Finally, Project Genesis,[8] the BISIH's first project focused on green finance, explored real-time tracking of environmental outcomes as well as smart-contract-based delivery of carbon credits in the context of tokenised green bonds.

These initiatives have generated valuable insights into climate risk data, analytical methodologies and persistent data gaps. Collectively, they have established a strong foundation of knowledge and technical expertise for further development.

In parallel with these efforts, the Bank of France, De Nederlandsche Bank and the Hong Kong Monetary Authority, as part of the BIS Innovation Network's Working Group on Green Finance, launched a project to address the challenge of incorporating physical climate related risks into financial decision-making through the use of digital twin technology. The initiative resulted in a proof of concept that developed a generic approach to creating digital twins capable of integrating external climate and environmental data with financial data to support forward-looking assessments of physical risks.

3. Project Gaia | Bank for International Settlements
4. Project Symbiosis | Bank for International Settlements
5. Project Ellipse | Bank for International Settlements
6. Project Viridis | Bank for International Settlements
7. Project NGFS Data Directory 2.0 | Bank for International Settlements
8. Project Genesis 1.0 | Bank for International Settlements

Framework implementation and scalability

Framework design

The Danu prototype is a reusable code base that can be further developed and adapted within an institution's own technical environment. A key objective was to explore how the digital twin concept could be applied in a manner that is practical, scalable and reusable to central banks and supervisory authorities. The framework was therefore designed based on the following requirements.

Integration of diverse data sources. Climate risk analysis typically requires information from multiple environmental, geospatial, asset and financial data sets from a wide range of sources and technical environments. To support this, the framework follows a medallion architecture that separates data ingestion, harmonisation and analytical layers, enabling consistent integration of information from heterogeneous sources (see Annex B).

Support for different user profiles. Central banks and supervisory authorities require both technical and non-technical users to interact with climate risk tools. The framework was therefore designed to support multiple user profiles, which are described in more detail in the following section.

Adaptability and scalability to institution-specific requirements. Although the framework was designed to be implemented across jurisdictions, institutions may have different analytical priorities, data availability and methodological requirements. While globally applicable methodologies can support broad reuse across jurisdictions, they may not fully capture local conditions. The framework therefore allows institutions to incorporate locally calibrated data sets and damage functions, enabling analyses to be adapted to their specific use case while maintaining a common analytical structure. In addition, its modular design enables the future integration of new data sources and analytical models as methodologies and analytical needs evolve.

Finally, the framework was designed to **minimise barriers to adoption** through the use of open technologies and openly available data sources, while allowing institutions to deploy and extend the prototype within their own environments. The open source technology stack and architecture are described in Annex B.

User roles

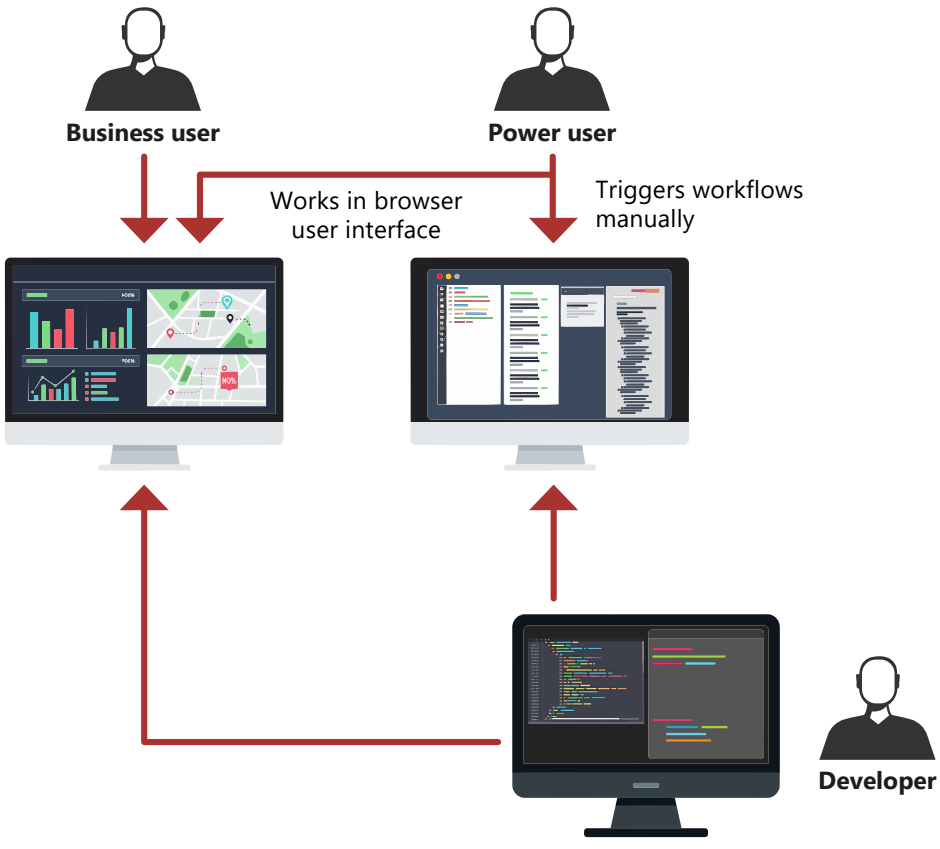
To ensure accessibility and usability across varying levels of technical expertise, the project distinguished between three main user groups:

- business users, who primarily interact with dashboards to explore outputs and run predefined scenarios;
- power users, who extend these capabilities through ad hoc analysis and experimentation with additional analytical features, such as incorporating new models or data sources; and
- developers, who maintain the system, manage the underlying data infrastructure and support the integration of new use cases.

To support users with different levels of technical expertise, the framework separates data and code management from analytical and visualisation components (Figure 2).

This architecture allows business users to focus on analytical outputs through dashboards, while enabling power users and developers to configure, extend and maintain the underlying framework. As a result, the prototype remains accessible to non-technical users while enabling advanced analysis and continuous system development.

User interactions with the framework — Figure 2



Data and methodology

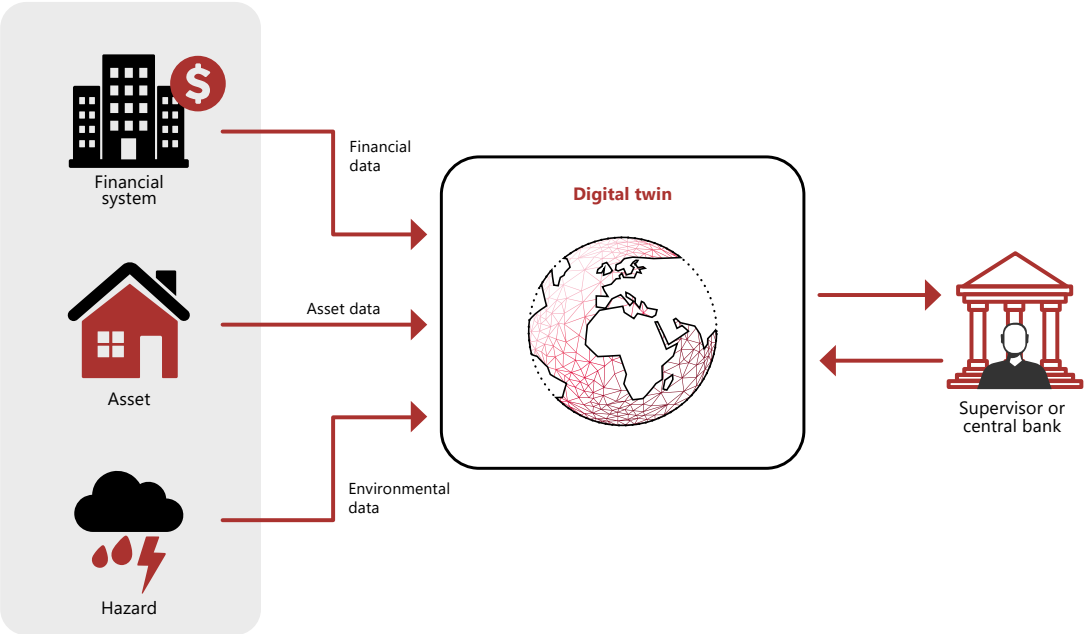
The prototype incorporates environmental and asset data sets that enable the monitoring of climate-related hazards and their impacts. In total, it draws on 17 environmental and asset data sources, ranging from satellite observations to global scenario map data sets (see Annex A for details). Figure 3 illustrates three categories of data required to translate a physical climate-related event into its potential impact on the financial system:

- **Environmental data provide information on the hazard.** Satellite-based sources are used to monitor current conditions in near-real time, while climate scenario data, such as those derived from Intergovernmental Panel on Climate Change (IPCC) projections, provide forward-looking insights.

- **Asset data capture the physical assets exposed to the hazard.** When combined with environmental data, they enable the estimation of potential damage, thereby informing vulnerability assessments. Asset data can be sourced from a variety of channels and are not limited to regulatory reporting by supervised entities. For the prototype, the focus is on real estate assets, using globally available open source data sets.

- **Financial data refer to institution-specific information associated with exposed assets.** Integrating these data enables the translation of physical damage into financial risk metrics. Users are expected to provide and integrate their own financial data; within the prototype, synthetic data were used for testing purposes.

Overview of the digital twin framework assessing the effects of physical climate risks on the financial system — Figure 3

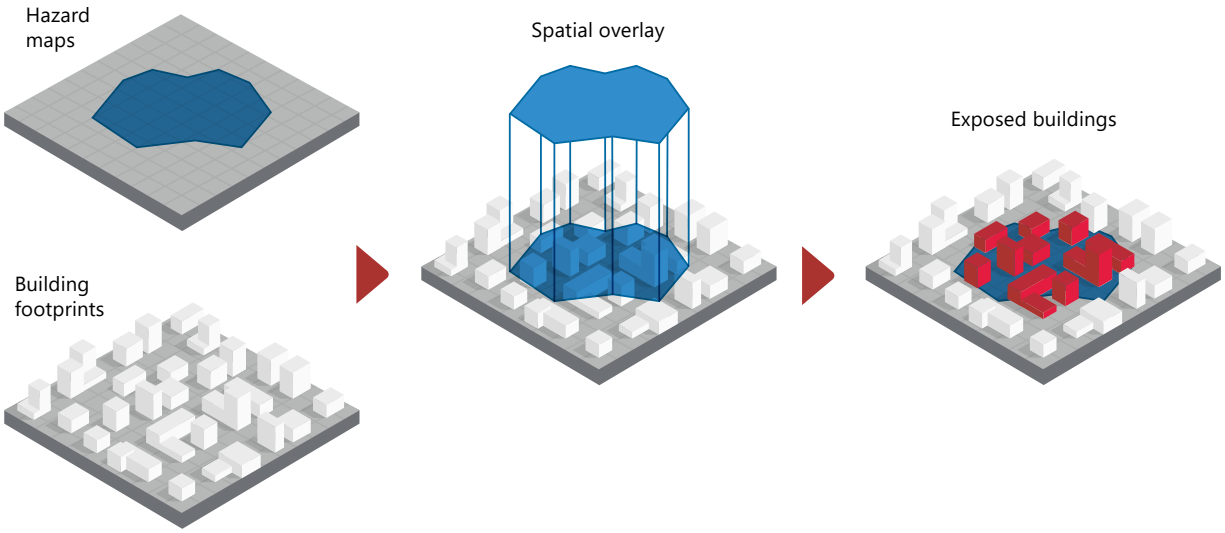


The methodology underlying the prototype implements the hazard exposure vulnerability finance (HEVF) model[9], commonly used in catastrophe modelling within the insurance industry. Rather than developing new hazard or damage models, the project focused on how established methodologies can be combined and operationalised within a common digital-twin-inspired framework for climate risk assessment.

The process begins with the hazard component, which defines the timing, location and intensity of an event. This is followed by the exposure component, where geospatial data are used to identify the assets located within the affected area and therefore exposed to the hazard. The vulnerability component then assesses the susceptibility of these exposed assets to damage under different hazard conditions.

Building on the methodology, the framework estimates how a climate-related event affects physical assets and translates these impacts into economic losses, such as repair costs and reductions in asset values. Where financial data are available, these estimates can be linked to information on loans and portfolios, helping users to understand how physical climate risks may affect financial institutions and financial stability.

Spatial overlay of hazard and asset data to identify exposed assets — Figure 4



9. A Dlugolecki (ed), Coping with climate change – risks and opportunities for insurers, Chartered Insurance Institute, 2009.

Figure 4 illustrates the core analytical process within the framework. Hazard data (eg on a flood, top left) are combined with geolocated asset data (eg building footprints, bottom left) through a spatial overlay process (centre) to identify assets exposed to a given hazard. The resulting layer (right) highlights potentially affected assets and provides the foundation for estimating physical damage, economic losses and, where financial data are available, financial impacts.

At the individual asset level, the damage is determined by combining hazard intensity with asset vulnerability. Asset vulnerability relates to the characteristics of assets, with parameters such as type, conditions or use. Using these parameters as inputs, a hazard-specific damage function estimates the physical damage to the asset.

For example, hazard intensity can correspond to flood level. Asset vulnerability can describe a property as a residential building of 120 square metres. The damage function will use these inputs to estimate the potential damage as a ratio of the property's value. For floods, damage ratios are derived from Huizinga et al. (2017)[10]. Further details are provided in Annex C.

Where financial data are available and linked to individual assets, estimated physical damage can be linked to financial exposures, enabling the assessment of potential financial impacts. This represents the finance layer of the HEVF framework. While synthetic data were used during development, the framework is designed to support the integration of real financial data.

The final stage of the methodology evaluates the implications of physical damage for financial institutions. Asset-level losses reduce the value of collateral securing loans and may increase borrower default risk. The current implementation supports the calculation of key metrics such as loss-given-default (LGD) and loan-to-value (LTV) ratios by comparing the value of outstanding loans with the value of collateral after the occurrence of a hazard event. This enables institutions to assess how physical climate risks may affect credit risk, portfolio losses, lending capacity and, ultimately, financial stability.

10. J Huizinga, H de Moel and W Szewczyk, "Global flood depth-damage functions: methodology and the database with guidelines", JRC Technical Reports, 2017.

Case study and results

This case study uses a real event to demonstrate how the digital twin concept can be applied to climate risk assessment by integrating hazard information with geospatial and asset data. This provides a basis for assessing the potential impacts of physical climate events and illustrates how the framework could be expanded through the incorporation of additional data sets, financial information and more advanced analytical capabilities.

Between July and August 2026, several wildfires occurred across Spain. The Danu framework was used to analyse these events using NASA FIRMS (Fire Information for Resource Management System) satellite data. The framework was able to identify exposed assets, estimate asset-level damage factors and translate these into indicative economic impacts. This provides a practical illustration of how the framework can support near real-time assessments of physical climate risks.

Map of damage intensity caused by wildfires in Spain on 26 July 2026 Figure 5

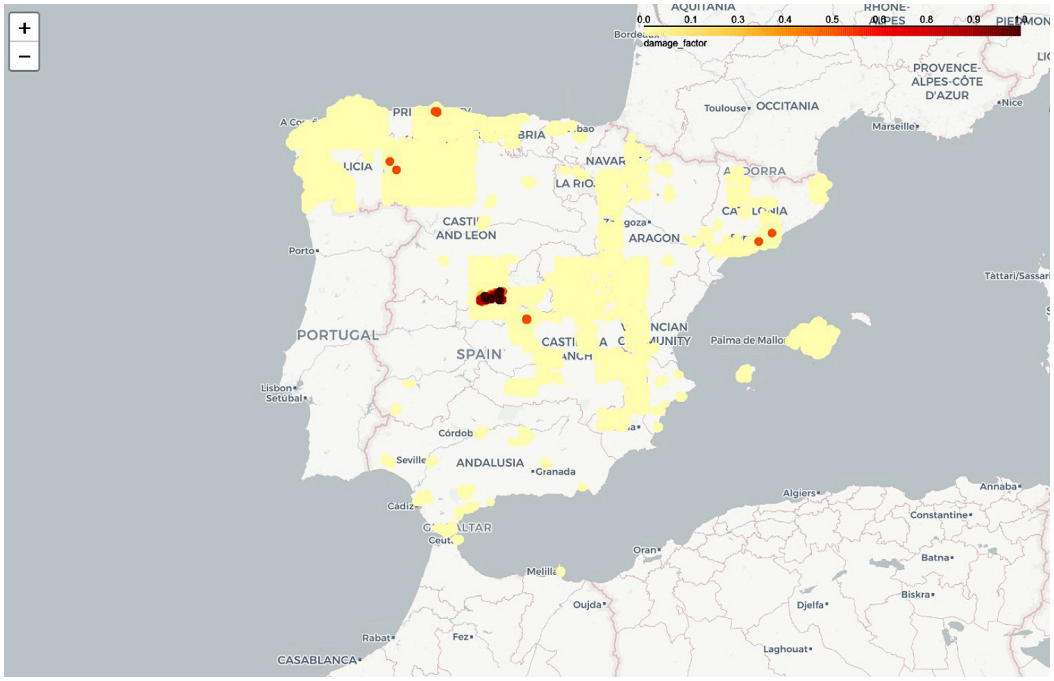


Figure 5 shows the estimated asset damage intensity resulting from the wildfire event in Spain on 26 July 2026. Danu calculates these estimates by combining physical damage information derived from satellite observations with geolocated asset exposure data. This allows the framework to identify which assets are affected by the wildfires and to estimate asset-level damage factors, ranging from 0 (no damage) to 1 (complete damage). Darker shades indicate higher levels of estimated damage.

The results show a concentration of high damage factors in parts of central Spain, where wildfire activity was most intense, while smaller clusters of affected assets are observed in other regions.

Danu user interface applied to Spain, capturing the wildfires on 26 July and 2 August 2026 — Figure 6

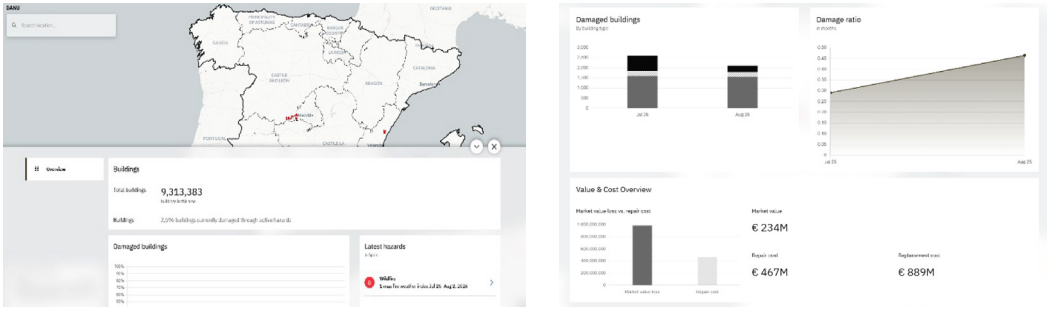


Figure 6 shows the indicative economic impact of the wildfires on 26 July and 2 August 2026 through the Danu user interface. Using satellite-derived wildfire observations and geolocated asset information, Danu identifies affected assets and estimates asset-level damage factors. These damage factors are combined with country-level assumptions on asset values and repair costs per square metre to estimate economic impacts.

The dashboard illustrates how the estimated impacts evolve over time. Between 26 July and 2 August 2026, the number of affected buildings decreased slightly, while the average damage ratio increased, indicating that although fewer assets remained affected, the estimated severity of damage among those assets increased. This demonstrates how the framework can support continuous monitoring of evolving hazard events and their impacts.

Based on the estimated damage factors and country-level valuation assumptions, the framework translates physical damage into indicative economic impacts. As shown in the dashboard, the estimated market value loss amounts to approximately €234 million, while repair costs are estimated at €467 million. The estimated replacement cost, representing the value required to fully replace the affected assets in the event of total destruction, is approximately €889 million.

These estimates are derived solely from environmental and asset data and do not require institution-specific financial information. They therefore provide an initial indication of the physical and economic consequences of the wildfire event, supporting situational awareness, impact assessment and risk monitoring.

Insights from the case study

The analysis shows that globally available environmental data can provide meaningful insights into hazard intensity, exposure and potential asset damage in near real time. This enables users to identify affected regions and assets as events unfold, estimate the extent of physical damage and quantify potential economic impacts through indicators such as repair costs, market value losses and maximum potential damage.

While the use case tested was an actual wildfire event, the framework would operate in a similar manner when used for scenario analysis. Rather than using observed hazard data, users can incorporate hazard projections derived from existing climate and hazard models, including data sets based on IPCC scenarios. These scenario inputs are available for all hazards currently within scope, enabling the assessment of potential future impacts within a common analytical framework.

The case study also showed that useful impact estimates can be generated using only publicly available environmental and asset data sets. This demonstrates that institutions can derive meaningful insights even where access to institution-specific data sets is limited. For central banks and supervisory authorities, such analyses can support the identification of vulnerable regions, while providing evidence to assess the potential economic consequences of climate-related events.

Where financial data are available, the analytical capabilities of the framework can be further extended. By linking asset-level damage estimates to institution-specific data sets, such as loan level or portfolio-level information, users can translate physical impacts into financial risk metrics and gain a more comprehensive understanding of the implications of climate-related hazards for the financial system.

Limitations and potential next steps

Limitations

The current implementation is subject to a number of limitations related to data availability, model assumptions and operational deployment.

The framework relies on external environmental, geospatial, asset and financial data sets, whose availability, quality and coverage may vary across hazards and jurisdictions. In addition, the hazard and damage models incorporate assumptions intended to provide a globally applicable baseline and may not fully reflect local conditions, building characteristics or market dynamics. Although the framework draws on established scientific methodologies and incorporates automated testing, certain components would benefit from further validation, calibration against local data sets and more comprehensive uncertainty analysis.

Institutions implementing the framework may therefore wish to supplement it with local data sets, modelling approaches and institution-specific assumptions where appropriate. Operational deployment would also require institution-specific arrangements for data governance, security, infrastructure, model management and ongoing maintenance.

Despite these limitations, the framework demonstrates how environmental, asset and financial information can be combined to support the assessment of climate-related physical risks. Future enhancements may incorporate additional data sets, refined methodologies and expanded analytical capabilities as data availability and modelling techniques continue to evolve.

Potential next steps

While the current prototype focuses primarily on monitoring and impact assessment, predictive and simulation capabilities represent a natural next step in the evolution of the framework towards a more comprehensive digital twin environment.

Initial testing was conducted in the area of flood forecasting, exploring how data-driven models could complement the existing framework. The project's predictive component was centred on supervised machine learning models that estimate short-term flood probability and spatial susceptibility using tabular features derived from topography, hydrometeorology and historical event footprints. The principal modelling paradigm employed gradient-boosted tree ensembles trained to predict event probability or probabilistic flood intensity at point locations. Labels for training were derived from historical event reconstructions and satellite proxies, with the prediction horizon tailored to operational needs (for example one-day-ahead to seven-days-ahead probability).

Future explorations of the framework could investigate increasingly sophisticated modelling of how hazards affect environmental, economic and financial systems more broadly. Potential approaches could be time series forecasting methods for hazards, physics-informed models to ensure consistency with established scientific relationships and machine learning techniques for asset-level risk assessment. Future research could explore how these approaches might be incorporated within the framework to provide increasingly more detailed and forward-looking assessments.

Summary

Project Danu explored how the digital twin concept can be applied within a climate risk and financial stability context in a manner that is practical, scalable and adaptable across central banks and supervisory authorities. The project demonstrated that environmental, geospatial, asset and financial data can be combined within a common analytical framework to monitor hazards and assess their potential impacts on physical assets, economic activity and, where relevant data are available, the financial system.

A key lesson from the project is that meaningful insights can be generated using globally available data sets. The case study showed how satellite observations and open source asset data can be used to identify affected assets, estimate damage and quantify indicative economic impacts during extreme weather events. This provides a foundation for monitoring risks even in situations where detailed institution-specific data are unavailable, supporting the assessment of potential implications for financial institutions and financial stability.

The project also generated practical insights into the application of the digital twin concept within a central banking and supervisory context. It focused on the interoperability across diverse technical environments, adaptability to institution-specific requirements, scalability as analytical needs evolve and low barriers to adoption. These requirements enabled the development of a reusable framework that can be deployed, extended and enhanced across a broad range of institutional contexts.

Project partners and acknowledgements

BIS Innovation Hub

Benjamin Bach, Developer

Darko Micic, Lead IT/Cloud Architect

BIS Innovation Hub Eurosystem Centre

Timothy Aerts, Adviser

Zooey Bossert, Adviser and Project Lead

Angela Dupont, Adviser

Christelle Therville, Project Portfolio Manager

Andras Valko, Acting Centre Head

BIS Innovation Hub Hong Kong Centre

Andrew Chau Chun Sing, Adviser

Bénédicte Nolens, Centre Head

Weeris Treeratanajaru, Adviser

Lucy Wong, Adviser

BIS Innovation Hub Singapore Centre

Geri Ang, Adviser

Carmen Arias, Adviser

Meena Chandra, Acting Centre Head

Patrick Hoffman, Adviser

Bank of France

Alexandre Le Douaron, Deputy Head of the Innovation Lab

Etienne de L'Estoile, Research Economist in the Financial Stability Directorate

Lisa Kerdelhue, Research Economist in the Financial Stability Directorate

Yann Marin, Deputy Director of the Financial Stability Directorate

Thierry Verdier, Tech Lead – Developer

De Nederlandsche Bank

Aimee Donsu, Data Scientist

Iman van Lelyveld, Department Head of Data Science Hub

Lisa Meijer, Data Scientist

Kitty Rang, Data Scientist

Hong Kong Monetary Authority

Renita Au, Senior Manager of Banking Conduct Division

Karen Li, Manager of Sustainable Banking Development Division

Warwick Yuen, Manager of Sustainable Banking Development Division

Ronald Young, Head of Sustainable Banking Development Division

The Monetary Authority of Singapore

Vivian Han, Assistant Director, Enterprise Knowledge Department

Jeremy Hor, Head, Enterprise Knowledge Department

David Nefzi, Assistant Director, Enterprise Knowledge Department

European Central Bank

Valérie M Saintot, Lead IT Innovation Expert, DG-IS Digital Office

Corina Wenzel, Senior IT Innovation Specialist, DG-IS Digital Office

Annex

A. Data sources

Data source	Provider	Data type	Purpose in framework	Hazard
OpenStreetMap (OSM)	OpenStreetMap/ Geofabrik	Asset data	Building footprints and asset locations for exposure mapping	Cross-hazard
Microsoft Global ML Building Footprints	Microsoft	Asset data	Building footprints for exposure assessment	Cross-hazard
Google Open Buildings	Google	Asset data	Building footprints for exposure assessment	Cross-hazard
geoBoundaries	geoBoundaries	Reference data	Administrative boundaries and geographic reference layers	Cross-hazard
Country building values	Danu reference data set	Reference data	Property valuation assumptions used to estimate damage and losses	Cross-hazard
OSM Buildings	Danu reference data set	Reference data	Mapping of building types and classifications	Cross-hazard
Copernicus Global Flood Monitoring (GFM)	Copernicus	Hazard data (observed)	Flood extent and flood detection	Flood
Sentinel-1 Ground Range Detected (GRD)	Copernicus	Hazard data (observed)	Satellite radar imagery supporting flood analysis	Flood
Copernicus GLO-30 Digital Elevation Model	Copernicus	Reference data	Elevation data supporting flood modelling	Flood
Aqueduct Floods	World Resources Institute (WRI)	Hazard data (scenario)	Riverine and coastal flood scenarios, depth maps, return periods	Flood
Global Forecast System (GFS)	National Oceanic and Atmospheric Administration (NOAA)	Hazard data (forecast)	Wind forecasts for cyclone modelling	Cyclone
International Best Track Archive for Climate Stewardship (IBTrACS)	NOAA National Centers for Environmental Information (NCEI)	Hazard data (historical)	Historical tropical cyclone tracks and characteristics	Cyclone

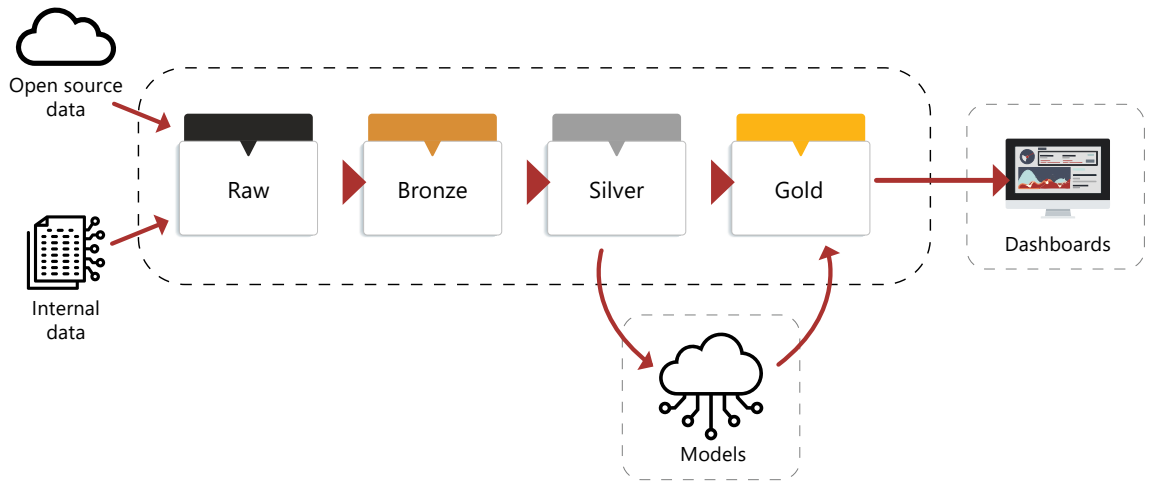
ERA5 reanalysis of winds	Copernicus Climate Data Store	Hazard data (reanalysis)	Wind field data for cyclone assessments	Cyclone
GDACS	Global Disaster Alert and Coordination System	Hazard data (event catalogue)	Disaster event monitoring and alerts	Cyclone
ESA WorldCover	European Space Agency (ESA)	Reference data	Land-cover information and exposure characterisation	Cyclone, drought
USGS PAGER (Prompt Assessment of Global Earthquakes for Response) reference data set	United States Geological Survey	Reference data	Construction characteristics and property weighting factors	Cyclone
STORM cyclone wind maps	STORM/Zenodo	Hazard data (scenario)	Future cyclone wind scenarios and return periods	Cyclone
NASA FIRMS MODIS (Moderate Resolution Imaging Spectroradiometer)	NASA	Hazard data (observed)	Active wildfire detection from satellite observations	Wildfire
NASA FIRMS VIIRS (Visible Infrared Imaging Radiometer Suite)	NASA	Hazard data (observed)	Active wildfire detection from satellite observations	Wildfire
ETH Fire Weather Index (FWI)	ETH Zurich	Hazard data (scenario)	Future wildfire risk and fire weather conditions	Wildfire
Drought indices (Standardised Precipitation-Evapotranspiration Index – SPEI)	Copernicus Climate Data Store	Hazard data (observed)	Drought severity indicators	Drought
ERA5-Land monthly	Copernicus Climate Data Store	Hazard data (reanalysis)	Temperature, precipitation and evapotranspiration variables	Drought

SoilGrids	ISRIC (International Soil Reference and Information Centre)	Reference data	Soil characteristics and clay content for vulnerability assessment	Drought
NEX-GDDP-CMIP6 (NASA Earth Exchange Global Daily Downscaled Projections)	NASA	Hazard data (scenario)	Climate projections used to derive future drought scenarios	Drought

B. High-level design

The Danu framework adopted a medallion architecture consisting of four layers: raw, bronze, silver and gold (Figure B.1). The raw layer maintains a one-to-one copy of data provided by external sources, preserving data sets in their native formats. Where required, only minimal processing is performed at this stage to enable subsequent ingestion. The bronze layer applies technical standardisation, including schema harmonisation, data-type assignment and conversion to common storage formats, providing a consistent technical foundation for further processing. The silver layer performs business-level transformation and enrichment, including data cleaning, spatial matching, integration across data sets and the derivation of intermediate analytical outputs. Finally, the gold layer produces reporting-ready data sets, including physical damage estimates, financial risk metrics and the aggregated outputs consumed by the application programming interface (API) and visualisation components.

High-level design of Danu — Figure B.1



The Danu prototype uses a collection of open source technologies that support data integration, workflow orchestration, analytics and visualisation. The core framework is implemented in Python and executes data-processing pipelines across the layers of the medallion architecture. Workflow scheduling and monitoring are provided by Apache Airflow, while DuckDB is used as the analytical engine for processing tabular and geospatial data sets. Persistent storage is provided through SeaweedFS, using an S3-compatible object storage interface, with JupyterLab supporting interactive analysis and framework development. The application layer consists of a Flask-based server that exposes APIs and analytical outputs, and a React-based user interface for visualisation and interaction. Together, these components provide a flexible, open and extensible technology stack that can be deployed and adapted within a variety of institutional environments.

C. Approach

Key methodological references

The framework draws on established methodologies from the climate risk, catastrophe modelling and insurance literature. These methodologies provide the conceptual foundation for the HEVF framework and the hazard-specific damage models implemented within the prototype. Additional scientific data sets and methodological approaches are drawn from organisations, including the European Commission's Joint Research Centre (JRC), the Copernicus Programme, NASA, ETH Zurich, Imperial College London and the United States Geological Survey (USGS) PAGER (Prompt Assessment of Global Earthquakes for Response) initiative. Detailed references for the individual methodologies are provided in the corresponding sections below.

Hazard-specific damage models

The framework employs hazard-specific damage functions to translate hazard intensity into estimates of physical damage. These functions are based on established scientific literature and catastrophe risk modelling methodologies and are designed to provide a consistent approach across hazard types.

Cyclone risk. Cyclone damage is modelled as a combined envelope capturing the maximum of wind-related and flood-related structural losses. Wind damage is estimated using sigmoidal vulnerability functions calibrated to local building typologies, based on the methodological framework established by Emanuel (2011)[11] and regionally parameterised by Eberenz et al (2021)[12]. The function determines the proportion of structural loss as a function of peak wind speed relative to a regional or building-specific half-damage threshold. Variations in regional building resilience and structural vulnerability are integrated using structural typology distributions derived from the USGS PAGER database.

11. K Emanuel, "Global warming effects on U.S. hurricane damage", Weather, Climate, and Society, vol 3, no 4, 2011, https://doi.org/10.1175/WCAS-D-11-00007.1.

12. S Eberenz, S Lüthi and D Bresch, "Regional tropical cyclone impact functions for globally consistent risk assessments", Natural Hazards and Earth System Sciences, vol 21, no 1, 2021, https://doi.org/10.5194/nhess-21-393-2021.

Drought risk. The drought damage model quantifies physical structural damage resulting from soil shrink-swell and drought-induced subsidence processes affecting buildings. The methodology combines drought severity, measured through the Standardised Precipitation-Evapotranspiration Index (SPEI), with clay content data sets from SoilGrids. The relationship between climate-driven drought anomalies, soil shrinkage and building vulnerability builds on empirical subsidence modelling frameworks (Barthelemy et al (2024))[13].

Flood risk. Flood damage is estimated using depth-damage functions derived from the JRC global flood damage database. The methodology links flood depth, property type and geographic region to a damage factor representing the proportion of asset value affected by a flood event. The underlying approach is based on the work of Huizinga et al (2017), who provide harmonised flood loss curves across multiple regions and asset classes.

Wildfire risk. For wildfire events, the current implementation assumes complete physical damage to assets located within detected wildfire areas. While intentionally conservative, this approach serves as a proof-of-concept methodology and demonstrates the integration of wildfire hazard data into the framework.
Future refinements may incorporate information on fire intensity, burn severity, proximity and asset-specific vulnerability characteristics to generate more granular damage estimates.

13. S Barthelemy, B Bonan, J-C Calvet, G Grandjean, D Moncoulon, D Kapsambelis and S Bernardie, "A new approach for drought index adjustment to clay-shrinkage-induced subsidence over France: advantages of the interactive leaf area index", Natural Hazards and Earth System Sciences, vol 24, no 3, 2024, https://doi.org/10.5194/nhess-24-999-2024.



Bank for International Settlements (BIS)

ISBN 978-92-9279-004-2