

H HEURISTICS

Navigating a Changing World

Polycrisis Resilience in an Era of Cascading Shocks

How Societies Can Adapt to Interconnected Global Risks

Why risks now travel together, how shocks cascade through coupled systems, and what governments, firms and communities can do to keep functioning — the evidence, a framework and a roadmap

H Heuristics · Hunter Hughes

Research Report · October 2026

© 2026 H Heuristics. All rights reserved.

Abstract

Major risks no longer arrive one at a time. The 2020s have delivered a pandemic, a war that disrupted energy and food markets, a surge in interest rates, record climate extremes and renewed geopolitical confrontation, each landing on societies whose reserves were drawn down by the last. This report examines how societies can remain functional, and keep improving, when shocks arrive together and propagate across the systems on which modern life depends.

The evidence base is stark. Seven of nine planetary boundaries are now transgressed, and the first Earth system tipping point, for warm-water coral reefs, has been crossed. The United Nations estimates the true annual cost of disasters at about US\$2.3 trillion, roughly ten times the direct losses officially recorded, and 80–90 per cent of catastrophe losses in emerging economies are uninsured. Public debt reached a record of about US\$102 trillion in 2024, and 61 developing countries now spend a tenth or more of government revenue on interest, eroding the buffers on which crisis response depends. Half of the experts surveyed by the World Economic Forum expect a turbulent or stormy world within two years.

The report argues that the unit of analysis for resilience must change from the single hazard to the interaction between systems. It reviews the causal pathways of entanglement (common stresses, domino effects and feedback loops), examines two real cascades, the 2025 Iberian blackout and the 2021 Texas freeze, and uses an illustrative network simulation to show how buffers, connectivity, simultaneous shocks and firebreaks determine whether a local failure stays local. In the simulation, raising shock tolerance from 18 to 30 per cent eliminates single-shock cascades; ten simultaneous failures can collapse a network that absorbs one; and a modular design with under 1 per cent of links crossing module boundaries confines a cascade to about 4 per cent of the system, a benefit that disappears once a few per cent of links cross.

Drawing on four cases, Bangladesh's cyclone preparedness, the Netherlands' Delta Programme, Finland's comprehensive security model and Jamaica's pre-arranged catastrophe finance, the report proposes a four-layer framework (prevent and decouple, absorb, adapt, transform) with seven design principles. It shows that reported returns on resilience investment run from about 2:1 to 10:1, sets out a playbook for six groups of actors, proposes a twelve-indicator dashboard, and closes with a three-phase roadmap, an account of the trade-offs and a statement of what evidence would change its argument.

Key findings

- **Entanglement, not hazard count, is the problem.** Crises are linked by common stresses, domino effects and feedback loops, so the combined harm exceeds the sum of the parts.
- **Buffers and firebreaks decide outcomes.** In an illustrative model, shock tolerance is the strongest lever, simultaneous shocks widen the danger zone, and genuine modularity contains cascades.
- **Resilience pays, if built early.** Reported benefit-cost ratios range from about 2:1 to 10:1; just 24 hours of warning can cut damage by about 30 per cent.
- **Adapted societies share a pattern:** early and sustained investment, institutions that outlast political cycles, distributed responsibility and continuous learning.
- **Pre-arranged finance and early warning are the highest-value starting points,** and both remain unevenly available: 60 per cent of countries have multi-hazard warning, and small island states only 43 per cent.

Keywords: polycrisis; resilience; cascading failures; systemic risk; tipping points; compound shocks; interdependent infrastructure; disaster risk finance; early warning; adaptation; polycentric governance; network models; sovereign debt; protection gap

Data and method: This report synthesises institutional and peer-reviewed research, including the World Economic Forum's *Global Risks Report 2026*; Lawrence et al. on the causal mechanisms of global polycrisis; the 2025 Planetary Health Check; the Global Tipping Points Report 2025; the WMO's *State of the Global Climate 2025*; UNDRR's *Global Assessment Report 2025*; Swiss Re's *sigma 1/2026*; UNCTAD's *A World of Debt*; the World Bank's *Lifelines*; the Global Commission on Adaptation; and primary and official sources on the Iberian blackout, Winter Storm Uri, Bangladesh, the Netherlands, Finland and Jamaica. Quantitative claims about the world are attributed inline. Figures 2, 3, 8 and 12 are conceptual schematics; Figure 7 is generated by an illustrative simulation described in the Appendix and is not an estimate for any real system; all other figures report values from the cited sources. The report is analytical rather than predictive.

Contents

1. Introduction: Resilience in a World of Entangled Risks	6
1.1 The problem: shocks that arrive together and travel together.....	6
1.2 Why the evidence now points one way.....	6
1.3 What this report argues	7
2. Understanding Polycrisis: Concepts and Mechanisms.....	9
2.1 Where the idea comes from	9
2.2 Three causal pathways	9
2.3 Why coupling matters more than shock size	10
2.4 A typology for practitioners.....	11
3. The Risk Landscape in 2026.....	12
3.1 The physical system: warming, boundaries and tipping points	12
3.2 The cost side: disasters are far larger than the official numbers.....	13
3.3 The buffer side: debt is eating the room to respond.....	13
3.4 The expert view: what the Global Risks Report 2026 says	14
3.5 Putting the evidence together.....	15
4. How Shocks Cascade: Lessons from Failures and Models	16
4.1 Two cascades, two lessons.....	16
4.2 What network models say	16
4.3 Seeing cascades coming.....	18
5. A Framework for Polycrisis Resilience	19
5.1 From bouncing back to bouncing forward.....	19
5.2 Four layers	19
5.3 Seven design principles.....	20
6. Lessons from Societies That Adapted	22
6.1 Bangladesh: from half a million deaths to a handful	22
6.2 The Netherlands: planning for uncertain futures	23
6.3 Finland: comprehensive security as a whole-of-society habit	23
6.4 Jamaica: pre-arranged finance meets a hurricane	23
6.5 What the cases have in common.....	24
7. The Economics of Resilience.....	25
7.1 Returns on resilience investment	25
7.2 Early warning: the best-value place to start.....	26
7.3 Paying for it: layers of risk finance.....	26
7.4 Why societies underinvest anyway	27
8. A Policy Playbook for Six Groups of Actors.....	28
8.1 National governments	28
8.2 Cities and subnational governments	28
8.3 Firms and infrastructure operators	28
8.4 The financial sector.....	28
8.5 Communities and civil society	29

8.6 The international system.....	29
9. Governance, Measurement and a Phased Roadmap	30
9.1 Governance: who owns the cross-cutting problem.....	30
9.2 Measurement: a short dashboard	30
9.3 A phased roadmap.....	31
10. Trade-offs, Risks and the Limits of This Analysis	32
10.1 Trade-offs.....	32
10.2 Ways the agenda can fail	32
10.3 Limits of the evidence.....	33
10.4 What would change this argument.....	33
11. Conclusion: Building the Capacity to Keep Going	34
Sources and Further Reading	35
Appendix: Methods, Data and the Illustrative Simulation	37
A.1 Approach and status of the figures.....	37
A.2 The cascade simulation	37
A.3 Reporting conventions	38

1. Introduction: Resilience in a World of Entangled Risks

The defining feature of the present decade is not any single crisis. It is that crises no longer stay in their lanes. Energy shocks become food shocks, food shocks become political shocks, and each one lands on societies whose buffers were spent on the last. This report asks what it takes for a society to keep functioning, and to keep improving, under those conditions.

1.1 The problem: shocks that arrive together and travel together

For most of the post-war period, risk management in governments and firms rested on a convenient assumption: that hazards could be assigned to separate categories, each with its own owner, its own budget and its own plan. A flood was a water ministry problem. A banking failure belonged to the central bank. A pandemic belonged to health authorities. The assumption was never entirely true, but it was close enough to organise institutions around.

It is far less true now. The 2020s have delivered a pandemic, a war in Europe that disrupted energy and grain markets, the sharpest rise in interest rates in a generation, a run of record-breaking climate extremes, and, in 2026, a Middle East conflict that has pushed energy prices up again. The World Bank's [Global Economic Prospects of June 2026](#) projects global growth of just 2.5 per cent this year, down from 2.9 per cent in 2025, and attributes the downgrade to higher energy prices, rising inflation and increased borrowing costs. No one of these events is unprecedented. What is new is the density of the sequence and the degree to which each shock has found the system already weakened.

Researchers and policymakers have started to use a single word for this condition: **polycrisis**. The historian Adam Tooze popularised the term in his [writing on the interplay of pandemic, war, energy and cost-of-living crises](#), tracing it to the French thinker Edgar Morin. The [Cascade Institute](#) and collaborators have since given it a more careful scientific definition, which section 2 develops. For now it is enough to say that a polycrisis is a cluster of crises in which the whole is worse than the sum of the parts, because the parts are causally entangled.

1.2 Why the evidence now points one way

Three bodies of evidence have converged over the past two years, and together they explain why the topic has moved from academic journals to the centre of policy agendas.

The first is the state of the Earth system. The [2025 Planetary Health Check](#) finds that seven of the nine planetary boundaries have now been transgressed, one more than in the previous assessment, with ocean acidification the latest addition. The [Global Tipping Points Report 2025](#) concludes that warm-water coral reefs have already passed a thermal tipping point. Figure 1 shows the current status of all nine boundaries.




Figure 1. Seven of nine planetary boundaries are beyond their safe limits.

H Heuristics, based on the [Planetary Health Check 2025](#) (Potsdam Institute for Climate Impact Research and partners). Binary status shown; the underlying assessment is quantitative for each boundary.

The second is the state of expert opinion. The World Economic Forum's [Global Risks Report 2026](#), published in January, ranks geoeconomic confrontation as the top risk both for 2026 and over the two-year horizon, and finds that half of surveyed experts anticipate a turbulent or stormy world over the next two years, up 14 percentage points from a year earlier. The [accompanying press release](#) notes that economic risks show the largest increases in ranking, while environmental risks have slipped in the short-term ranking even as they dominate the ten-year outlook.

The third is the state of the balance sheet. The United Nations Office for Disaster Risk Reduction's [Global Assessment Report 2025](#) estimates that the true annual cost of disasters, once cascading and ecosystem effects are counted, is about US\$2.3 trillion, roughly ten times the direct losses recorded in official statistics. UNCTAD's [A World of Debt](#) reports record public debt and a record number of developing countries spending a tenth or more of government revenue on interest. Societies are being hit harder and, at the same time, have less room to absorb the blow.

1.3 What this report argues

This report makes four arguments, and the rest of the document defends them in turn.

- **The unit of analysis must change.** Resilience planning organised around single hazards will systematically underestimate risk, because the dominant losses come from interactions: common stresses, domino effects and feedback loops between systems.
- **Connectivity is a double-edged design variable, not a fact of life.** Network science shows that the same links that make systems efficient can carry failure across them. An illustrative simulation developed for this report shows that cascades are driven by the interaction of coupling and buffers, and that deliberate firebreaks can contain them.

- **Resilience is cheaper than recovery, but only when it is built before the shock.** The best available estimates put returns on resilience investment at several dollars per dollar spent, yet pre-arranged finance, early warning and protection for lifelines remain underfunded.
- **Societies that have adapted share a recognisable pattern.** Bangladesh's cyclone preparedness, the Netherlands' Delta Programme, Finland's comprehensive security model and Jamaica's pre-arranged catastrophe finance differ in context, but all combine early investment, distributed responsibility and learning loops.

The report offers a framework built on four layers (prevent and decouple, absorb, adapt, transform), a portfolio of actions for six groups of actors, a phased roadmap and a short list of indicators. It is analytical rather than predictive: it does not forecast which crisis comes next, but argues that societies can reduce the damage from whichever does.

Table 1. The evidence base in brief: five indicators of a more entangled risk environment

Indicator	Latest reading	Source
Planetary boundaries transgressed	7 of 9 (up from 6 of 9 previously)	Planetary Health Check 2025
Global warming, 2025 (vs 1850–1900)	About 1.43°C; second or third warmest year	WMO State of the Global Climate 2025
True annual cost of disasters	About US\$2.3 trillion, vs about US\$202 billion direct	UNDRR GAR 2025
Global public debt, 2024	US\$102 trillion, of which about US\$31 trillion in developing countries	UNCTAD, A World of Debt
Experts expecting a turbulent or stormy world within two years	About half (up 14 points on 2025)	WEF Global Risks Report 2026

Figures are as reported by the cited sources; see section 3 for caveats on comparability.

2. Understanding Polycrisis: Concepts and Mechanisms

A polycrisis is not simply a bad year with many headlines. It is a specific condition in which several systems are pushed out of equilibrium together and in which the links between them, not just the shocks themselves, do the damage. The distinction matters because it changes what resilience requires.


2.1 Where the idea comes from

The word has a longer history than its recent popularity suggests. The French complexity theorist Edgar Morin used it to describe the interlocking problems of planetary development. Jean-Claude Juncker revived it in 2016 to describe the overlapping crises facing the European Union. Tooze's 2022 column in the *Financial Times* made it a staple of public debate, defining it by the idea that disparate shocks interact so that the whole is worse than the sum of the parts.

The most rigorous treatment to date is [Lawrence, Homer-Dixon, Janzwood, Rockström, Renn and Donges \(2024\)](#) in *Global Sustainability*, summarised on the [Stockholm Resilience Centre's publication page](#). The authors define a global crisis as arising when one or more fast-moving trigger events combine with slow-moving stresses to push a global system out of its established equilibrium and into a volatile and harmful state. A polycrisis exists when several such crises are causally entangled, so that the combined harm to humanity's prospects substantially exceeds what the separate crises would produce on their own.

2.2 Three causal pathways

The same paper identifies three pathways by which crises become entangled. They are worth stating plainly because each calls for a different kind of response. Figure 2 summarises them.



Fast-moving trigger + slow-moving stress → disequilibrium across several systems at once

Figure 2. Three causal pathways connect crises: common stresses, domino effects and feedback loops.

H Heuristics, adapting the framework of [Lawrence et al. \(2024\)](#). Examples are illustrative.

Common stresses. A single underlying driver acts on many systems at once. Climate heating raises cooling demand on power grids, stresses crops and water supplies, and increases disease risk simultaneously. Ageing populations strain health systems, pension systems and labour markets at the same time. Because the driver is shared, the systems it affects tend to fail in the same years, which defeats diversification strategies that assume independent failures.

Domino effects. A failure in one system becomes a shock to another. A power outage stops water pumps, hospital equipment, payment terminals and mobile networks, and then the supply chains that depend on them. These are the cascades most familiar from infrastructure failures, and the ones most amenable to engineering and design responses.

Inter-systemic feedbacks. Effects loop back to reinforce the original crisis. Sovereign debt stress forces cuts in adaptation and health spending, which raises disaster and disease losses, which raises debt stress. In the climate system, thawing permafrost releases carbon that accelerates warming. Feedbacks are the most dangerous pathway because they can sustain a crisis after the original trigger has passed.

2.3 Why coupling matters more than shock size

The central insight of the complexity literature is that the vulnerability of a system depends on its structure as much as on the size of the forces acting on it. Figure 3 shows a stylised map of eight systems on which modern societies depend. Their connections run in many directions, and some, such as energy to digital networks or water to food, are strong enough that failure in one makes failure in the other close to certain.




Figure 3. A stylised map of coupled societal systems. Line weight indicates the strength of dependency.

H Heuristics. Conceptual schematic: link strengths are qualitative judgements for illustration, not measurements.

Two features of this structure deserve emphasis. First, the map is dense: most systems are connected to most others, directly or through one intermediary, so a shock that begins anywhere has many routes to travel. Second, the climate and ecosystem node is not just another system but a source of common stress for all of them, which is why climate change figures in nearly every account of polycrisis.

The research literature on networks has made a related point with unusual sharpness. Buldyrev and colleagues showed in [Nature \(2010\)](#), using the 2003 Italian blackout as a case, that interdependent networks can be dramatically more fragile than isolated ones, because a failure in one network can trigger failures in the other that then feed back. Dirk Helbing's [analysis of globally networked risks](#) made the policy implication explicit: connectivity that improves day-to-day efficiency can also remove the natural firebreaks that once limited the spread of failure.

2.4 A typology for practitioners

Table 2 translates the three pathways into a practical typology. The last column matters most for design: each pathway implies a different class of resilience measure, which is why a portfolio of measures, not a single flagship programme, is required.

Table 2. Pathways of crisis entanglement and the response each implies

Pathway	Typical example	Early signal	Primary response
Common stress	Heat wave raising power demand, cutting crop yields and straining hospitals in the same week	Rising correlation among sector losses in the same period	Reduce the shared driver; diversify; adapt every sector to the same hazard
Domino effect	Grid failure disabling water, telecoms, payments and logistics	Dependency maps showing single points of failure	Redundancy; modular design; backup for critical nodes; firebreaks
Feedback loop	Debt stress cutting adaptation spending, raising disaster losses and debt	Same indicator worsening in successive years with no external trigger	Break the loop: pre-arranged finance, debt-pause clauses, counter-cyclical buffers

H Heuristics synthesis of Lawrence et al. (2024) and the sources cited in sections 3 and 4.

3. The Risk Landscape in 2026

The evidence that societies are being pushed harder and have less room to absorb shocks is now substantial. This section reviews it under three headings: the physical system, the economic buffers and the expert view of what lies ahead.

3.1 The physical system: warming, boundaries and tipping points

The WMO's [State of the Global Climate 2025](#) puts 2025 at about 1.43°C above the 1850–1900 baseline, the second or third warmest year on record depending on the dataset, and confirms that 2015–2025 were the eleven warmest years in the record. The same assessment treats 2024, at roughly 1.55°C, as the warmest year on record; the [WMO's provisional release](#) set out the early picture. The 1.5°C threshold in the Paris Agreement is a long-term average rather than a single-year limit, but single years above it show how little margin is left.

The significance for resilience lies in thresholds. The [2022 reassessment by Armstrong McKay and colleagues](#) concluded that five of sixteen climate tipping elements, including parts of the Greenland and West Antarctic ice sheets and abrupt permafrost thaw, may already be triggerable at current temperatures, and that further elements become possible near 1.5°C. The [Global Tipping Points Report 2025](#), which draws on more than 160 scientists at 87 institutions in 23 countries, goes a step further: it concludes that the first Earth system tipping point has been crossed, with the central estimate of the warm-water coral reef threshold at 1.2°C, and warns that Amazon dieback and the weakening of ocean circulation are next in line. [Time's coverage](#) summarises the finding for a general audience. The IPCC's [Sixth Assessment Report](#) had earlier judged the risk of triggering tipping points to be high by around 2°C and very high at 2.5–4°C.




Figure 4. Observed warming already sits inside the range where some tipping thresholds lie.

H Heuristics, using [WMO](#) (observed warming), [Global Tipping Points Report 2025](#) (coral threshold) and [IPCC AR6](#) (risk levels). Single years are not the same as the long-term averages to which thresholds refer.

Two points follow for resilience planning. Tipping elements interact, so a threshold crossed in one part of the system raises the chance of crossing another. And some changes are irreversible on any planning timescale, which means that adaptation to them cannot be postponed until the evidence is complete. The

IPCC estimates that between 3.3 and 3.6 billion people live in contexts that are highly vulnerable to climate change, which is the human exposure on which these physical risks fall.

3.2 The cost side: disasters are far larger than the official numbers

The accounting of disaster losses has long understated the problem. The [UNDRR Global Assessment Report 2025, Resilience Pays](#), reports that direct disaster costs averaged US\$70–80 billion a year between 1970 and 2000, rose to US\$180–200 billion a year between 2001 and 2020, and now run at about US\$202 billion annually. When the cascading effects on health, supply chains, finance and ecosystems are counted, the true figure is about US\$2.3 trillion a year, nearly ten times the official direct number. The [PreventionWeb summary](#) of the report adds that the burden, and the debt it creates, falls disproportionately on developing countries.

Insurance data point the same way. The Swiss Re Institute's [sigma 1/2026](#) puts insured natural catastrophe losses in 2025 at about US\$107 billion and total economic losses at about US\$220 billion, with roughly 49 per cent insured, the highest share in sigma's records. Swiss Re's [press release](#) reports that a record 92 per cent of insured losses came from secondary perils such as wildfire, flood and severe convective storms, that the Los Angeles wildfires alone generated about US\$40 billion of insured losses, and that more than 80 per cent of the uptrend in weather-related insured losses is explained by growth in exposure. Crucially, protection gaps in emerging economies are wide: 80–90 per cent of catastrophe losses there are typically uninsured. The 2025 total was below the long-term trend Swiss Re implies, which the institute attributes to favourable variability rather than any easing of underlying risk.




Figure 5. Disaster costs are much larger than counted, and most of the loss in emerging economies is uninsured.

H Heuristics, using [UNDRR GAR 2025](#) and [Swiss Re sigma 1/2026](#). Panel A shows the midpoints of the reported ranges for the two historical periods; “true cost” includes cascading and ecosystem costs. In panel B the emerging-economy bar is drawn at the midpoint of the reported range.

3.3 The buffer side: debt is eating the room to respond

A society's capacity to absorb a shock depends on what it has in reserve, and public balance sheets are the most important reserve of all. Here the picture is troubling. UNCTAD's [A World of Debt](#) reports that global public debt reached a record of about US\$102 trillion in 2024, with developing countries accounting for about US\$31 trillion and their debt growing roughly twice as fast as that of advanced economies since 2010. A record 61 developing countries now allocate 10 per cent or more of government revenue to interest. Secondary reporting of the 2025 data puts the net interest bill of developing countries at close to US\$1

trillion, nearly three times the US\$363 billion of 2010, and notes that in countries home to roughly 3.3 billion people, governments spend more on interest than on health or education.




Figure 6. Debt service is crowding out the fiscal buffers on which crisis response depends.

H Heuristics, using [UNCTAD, A World of Debt](#). The 2010 and 2025 net interest figures are as reported in secondary coverage of the UNCTAD data; the 2025 figure is rounded.

The connection to polycrisis is direct. A government that spends a tenth of its revenue on interest has less to spend on early warning, flood defences, social protection and health systems, which are precisely the investments that limit the damage of the next shock. That is a feedback loop of the third kind described in section 2: debt stress reduces resilience spending, which raises losses, which raises debt. Section 7 returns to how pre-arranged finance can break it.

3.4 The expert view: what the Global Risks Report 2026 says

The World Economic Forum's [Global Risks Report 2026](#) offers a useful cross-check from a different angle: the judgement of several hundred risk experts rather than the physical or financial record. Table 3 summarises the findings that bear most directly on resilience planning.

Table 3. What the Global Risks Report 2026 reports

Finding	Detail	Implication for resilience
Top risk, 2026 and two-year horizon	Geoeconomic confrontation, up eight positions on the two-year view versus last year	Supply-chain and trade dependencies are now a first-order resilience issue
Mood of respondents	About half expect a turbulent or stormy world within two years, up 14 percentage points	Planning assumptions of stability are no longer credible
Economic risks	Largest rises in ranking over two years: downturn, inflation and asset-bubble burst all climb	Fiscal and financial buffers matter as much as physical protection
Environmental risks	Extreme weather falls from 2nd to 4th in the short-term ranking, while still dominating the long term	Short-term attention is shifting; long-term exposure is not

Source: [World Economic Forum](#); H Heuristics interpretation in the final column. The full two-year ranking is in the report itself.

Artificial intelligence is the notable outlier: adverse outcomes rank 30th over two years but 5th over ten, a slow-moving risk that needs governance before it becomes acute. One caution: expert rankings are sensitive to recent events and to survey framing, and a fall in rank for extreme weather does not mean the risk has

diminished, as sections 3.1 and 3.2 show. The survey's value is as a picture of how decision-makers perceive the landscape, which shapes where resilience money is spent.

3.5 Putting the evidence together

Read together, the four bodies of evidence in this section support five conclusions that shape the rest of the report.

- **The physical stress is rising and partly locked in.** Warming is already within the range of some tipping thresholds, so adaptation cannot wait for mitigation to succeed.
- **Losses are larger than they look.** Official figures capture only about a tenth of the true cost of disasters, so resilience investments appear less valuable than they are.
- **The shock absorbers are thin where they matter most.** Debt service and uninsured losses are concentrated in the countries with the greatest exposure.
- **Decision-makers see the landscape changing.** Expert opinion has moved towards expecting turbulence, and towards economic and geopolitical risks that interact with environmental ones.
- **None of this is deterministic.** Each indicator reflects choices about buffers, standards and finance that can be changed, which is the subject of the sections that follow.

4. How Shocks Cascade: Lessons from Failures and Models

Cascades are not mysterious. Real failures leave records, and network science offers simple models that reproduce their main features. Putting the two together shows which design choices make a society fragile to compound shocks and which make it robust.

4.1 Two cascades, two lessons

The Iberian blackout, April 2025. On 28 April 2025 a blackout struck the whole of the Iberian Peninsula. The expert panel convened by ENTSO-E, the European association of transmission system operators, published its final report on 20 March 2026. As summarised by [pv magazine](#), the panel treated a surge in voltage as the immediate cause but stressed that it was not the only one. Conventional plants did not supply the voltage control the system operator expected under existing rules; some network components were operated manually; protection settings that disconnect equipment during voltage surges did not match requirements; many renewable plants operated at a fixed power factor, which limited their response; and oscillations and rapid generation losses combined to drive voltages up and trigger cascading disconnections. [T&D World](#) described it as a unique and perplexing event, a reminder that even well-understood systems can fail in unfamiliar ways.

The lesson is not about any particular technology, and the panel's chair was explicit that the issue was voltage control rather than renewable generation as such. The lesson is that a highly engineered system failed through the interaction of several modest weaknesses, none of which was individually decisive, and that the system had little capacity to contain the failure once it started. Reports differ on the details of restoration times and on some numbers, so this report relies only on the qualitative finding of multiple interacting causes.

Winter Storm Uri, Texas, February 2021. A prolonged freeze in February 2021 exposed the interdependence of gas and electricity. The joint federal inquiry by FERC and NERC, summarised in [T&D World](#), identified freezing of generation equipment and shortfalls in natural gas fuel supply as the two largest causes. The grid operator ordered about 20,000 MW of rolling blackouts, the largest manually controlled load shed in US history; more than 4.5 million people lost power, some for as long as four days; and estimates of the death toll range from just over 200 to as many as 700.

Uri is a textbook domino cascade. Cold disabled gas wells and pipelines, which removed fuel from power plants, which removed power from the compressors and electrically driven equipment that gas systems rely on. The regulatory response, including [new cold-weather reliability standards approved by FERC two years later](#), focused on weatherisation of both power and gas assets. The point for resilience planning is that the vulnerability lay in the coupling between two systems that had been regulated and engineered separately.

4.2 What network models say

Network science provides a way to ask these questions in general. In the threshold models pioneered by [Duncan Watts](#), each system in a network fails when the share of its neighbours that have failed passes a tolerance threshold. Such models reproduce a striking feature of real cascades: a system can be robust to most disturbances yet fragile to a rare few. [Haldane and May's analysis of banking ecosystems](#) made the same point for finance, showing that interconnection can absorb small shocks while amplifying large ones.

To make the logic concrete, this report uses a version of the Watts model, run on 1,000 coupled systems. It is an illustrative simulation, not an estimate for any real system, and its structure and parameters are set out in the appendix. Figure 7 shows three results.




Figure 7. An illustrative cascade model: connectivity, buffers, compound shocks and firebreaks.

H Heuristics illustrative simulation (Watts-type threshold model; 1,000 systems; 60–80 runs per point; see Appendix). Not calibrated to any real system. Panel A: random network, single initial failure. Panel B: tolerance 18 per cent, one versus ten simultaneous initial failures. Panel C: 20 modules, average of 4.5 links per system, tolerance 18 per cent.

Buffers are the strongest lever (panel A). With a tolerance of 30 per cent, meaning that a system survives unless nearly a third of its neighbours fail, a single shock almost never cascades at any level of connectivity: under 1 per cent of systems fail. With a tolerance of 18 per cent, a single shock takes down on average about 87 per cent of the network at the worst point, around 4 to 5 links per system. With a tolerance of 12 per cent, near-total cascades occur across almost the whole range from 3 to 8 links.

Connectivity is not monotonic (panel A). At very low connectivity, failures have nowhere to go. At intermediate connectivity, they spread freely. At high connectivity, in this model, each neighbour's failure is a small share of a well-connected system's load, so cascades die out; for the 18 per cent case, a single shock fails about 2 per cent of the network at 7 links and almost none beyond. This does not mean that dense networks are safe in reality. The model has no mechanism for the correlated, large shocks that Haldane and May show can overwhelm densely connected finance. The point is narrower: the danger zone sits at intermediate connectivity, which is where many real systems operate.

Compound shocks widen the danger zone (panel B). When ten systems fail at once instead of one, the range of connectivity over which near-total collapse occurs stretches substantially. At 18 per cent tolerance, a network with 7 links per system, which absorbs a single shock almost completely, fails nearly entirely under ten simultaneous failures (98 per cent); at 8 links the figure is 68 per cent, and at 9 links it is 30 per cent. This is the stylised version of the polycrisis problem: a design that is robust to any one shock can be fragile to several arriving together.

Firebreaks work (panel C). When the same 1,000 systems are arranged in 20 modules with sparse links between them, the share of links that cross module boundaries determines the outcome. With no cross-module links, a single failure takes down only about 4 per cent of the network, which is the extent of one module. With 1 per cent of links crossing boundaries, the figure is about 12 per cent. At 2 per cent it is about 57 per cent, and at 3 per cent it is about 79 per cent, close to the 82–88 per cent seen at 5 per cent or more. The transition is sharp: a few cross-links are enough to make a modular network behave like an unstructured one, so firebreaks must be genuine rather than nominal.

4.3 Seeing cascades coming

A final lesson from the science concerns warning. [Scheffer and colleagues](#) showed that many complex systems display generic early-warning signals before a critical transition: they recover more slowly from small disturbances, and their fluctuations grow and become more correlated. These signals have been demonstrated in ecological, climatic and financial settings. They are not reliable enough to time a crisis, but they are useful as a prompt for investigation, and they argue for monitoring the recovery speed of critical systems after minor disruptions, which is cheap and rarely done.

Table 4. Design lessons from the cascades and the model

Observation	Evidence	Design implication
Failures come from interacting weaknesses, not single faults	Iberian blackout: voltage control, protection settings, oscillations, manual operation	Audit interactions between components, not only components; test combinations
Coupling between separately managed systems is a hidden vulnerability	Texas 2021: gas supply and electricity	Joint planning and shared standards for interdependent infrastructures
Buffers dominate outcomes	Simulation panel A: tolerance 30 per cent virtually eliminates cascades	Fund slack, reserves and redundancy as a primary strategy, not waste
Compound shocks break designs tested one at a time	Simulation panel B: ten simultaneous failures	Stress-test for simultaneous and sequential events, not single scenarios
Firebreaks must be real	Simulation panel C: sharp transition above a few cross-links	Genuine modular boundaries, islanding capability, tested isolation

H Heuristics synthesis. Simulation findings are illustrative; see Appendix.

5. A Framework for Polycrisis Resilience

If cascades are driven by coupling, buffers and the absence of firebreaks, the response must act on each. This section sets out a four-layer framework and seven design principles that follow from the evidence in sections 3 and 4.

5.1 From bouncing back to bouncing forward

The word resilience is used in many ways, and the confusion has real costs. In engineering, it often means the speed with which a system returns to its prior state. In ecology and the social sciences, the concept has a richer lineage. C. S. Holling's [1973 paper on the resilience and stability of ecological systems](#) defined resilience as the capacity of a system to absorb disturbance and still persist. [Walker, Holling, Carpenter and Kinzig \(2004\)](#) extended the idea to social-ecological systems and distinguished resilience from two related capacities: adaptability, the ability of actors to manage the system's response, and transformability, the capacity to create a fundamentally new system when the existing one becomes untenable.

[Folke and colleagues \(2010\)](#) built these ideas into what is now called resilience thinking. The [Stockholm Resilience Centre](#) stresses that this is richer than bouncing back: it is about navigating the future in relation to thresholds and alternative development pathways. That emphasis matters for polycrisis. When tipping points have been crossed or debt has erased a country's fiscal space, returning to the earlier state may be neither possible nor desirable. A resilience strategy that cannot contemplate transformation is a strategy for the last crisis.

5.2 Four layers

The framework used in this report has four layers, shown in Figure 8. They are ordered by how much change they imply, and a society needs all four, because each handles a different kind of shock.

Prevent and decouple. Reduce the likelihood and size of shocks where this is possible, limit exposure to them, and build firebreaks between systems so that failures do not propagate. This is the layer that acts directly on the coupling identified in section 4.

Absorb. Hold enough slack that a shock can be taken without collapse: reserves, redundancy, insurance and pre-arranged finance. This is the layer that acts on the tolerance threshold in the simulation and on the buffer erosion documented in section 3.

Adapt. Change how systems operate so that they keep their function under altered conditions: early warning systems, resilient infrastructure, flexible planning, adaptive social protection. Adaptation is the dominant strategy against slow-moving stresses.

Transform. Change the structure of a system when its present form cannot be sustained: retreat from locations that cannot be defended, new models of energy, food and finance, and institutional redesign. Transformation is the hardest layer politically and the one most often deferred until a crisis forces it.




Figure 8. The H Heuristics framework: four layers of resilience and five enabling conditions.

H Heuristics, building on Holling (1973), Walker et al. (2004) and Folke et al. (2010). Conceptual framework.

5.3 Seven design principles

Seven principles run across the layers. They are drawn from the evidence reviewed here and from the case studies in section 6.

- **Map dependencies and find the critical nodes.** Most societies do not know which few facilities, firms or links carry a disproportionate share of systemic risk. A dependency map, built across sectors rather than within them, is the foundation of everything else.
- **Build real firebreaks.** Modular design, islanding capability and tested isolation procedures contain cascades. The simulation shows that the boundaries must be genuine: a few cross-links erase the benefit.
- **Keep slack.** Buffers and redundancy look wasteful in calm periods and are the single most effective protection in a crisis. Slack should be explicit in budgets and standards rather than squeezed out as inefficiency.
- **Diversify, but check for common stress.** Diversification protects against independent failures, but common stresses such as heat or drought hit many sources at once. Diversification must be tested against shared drivers.
- **Pre-commit.** Decisions, money and authority arranged before a shock are far faster than those negotiated during one. This covers pre-arranged finance, standing emergency powers with sunset clauses, and agreed triggers for action.
- **Distribute authority and build trust.** Resilience in complex systems comes from many decision-makers with local knowledge and clear responsibilities, coordinated through shared awareness rather than central command. [Elinor Ostrom's work on polycentric governance](#) is the best-developed account of why this works, and why it depends on trust.

- **Learn continuously.** Drills, stress tests, after-action reviews and monitoring of recovery speed convert each shock into improved capacity. Societies that do not learn pay the same price repeatedly.

Table 5. The seven principles and what happens without them

Principle	In practice	Failure when absent
Map dependencies	Cross-sector dependency register; named critical nodes	Hidden single points of failure (cf. Texas 2021)
Build firebreaks	Modular grids, islanding, supplier diversity at critical nodes	Local fault becomes system-wide cascade (cf. Iberia 2025)
Keep slack	Reserves, stockpiles, spare capacity, fiscal buffers	Low tolerance threshold: small shocks cascade
Diversify, test for common stress	Multiple sources and routes, checked against shared hazards	Apparent diversity fails together in a heat or drought year
Pre-commit	Pre-arranged finance, triggers, delegated emergency authority	Weeks lost negotiating while losses accumulate
Distribute authority, build trust	Local responsibility, networks, shared situational awareness	Central bottlenecks; compliance collapses when trust is low
Learn continuously	Drills, stress tests, public after-action reports	Same failure repeated; recovery speed unmeasured

H Heuristics synthesis.

6. Lessons from Societies That Adapted

Theory is useful only if practice bears it out. This section examines four cases in which societies built resilience to specific, serious hazards and achieved measurable results. None is a perfect model, and each has limits, but together they show what the principles look like when they are made to work.

6.1 Bangladesh: from half a million deaths to a handful

Bangladesh is among the most cyclone-exposed countries in the world and one of the poorest of those exposed. The 1970 Bhola cyclone killed hundreds of thousands of people; a World Meteorological Organization review put the toll at about 300,000, while other estimates are higher still. When Cyclone Sidr struck in 2007, with winds of around 260 km/h, the death toll was about 4,200 according to most sources. When Cyclone Amphan struck in 2020, the reported toll in Bangladesh was under 30. [Dialogue Earth](#) and the [Business Standard](#) document the figures and note that estimates for the historical events vary widely.




Figure 9. Bangladesh: cyclone deaths have fallen by orders of magnitude as shelters and volunteer networks expanded. H Heuristics, using [Dialogue Earth](#) and [Business Standard](#). Death tolls are contested and shown as approximate, on a logarithmic scale; the Bhola figure is the WMO review estimate. Shelter figures are approximate.

The mechanisms are well documented. Two years after Bhola, the Bangladesh Red Crescent Society, working with the government, launched the Cyclone Preparedness Programme, which today relies on about 76,000 volunteers, half of them women, to carry warnings to households and organise evacuation. In 1970 the country had only 44 cyclone shelters; it now has about 5,500, including some 1,500 multi-purpose centres. Forecasting, embankments and coastal afforestation complemented the shelters. In terms of the principles above, Bangladesh combined pre-commitment (standing warning and evacuation procedures), distributed authority (a volunteer network embedded in villages), and learning (each cyclone prompted improvements).

The limits are equally instructive. Commentators note that losses of livelihoods, crops, livestock and infrastructure have fallen far less than deaths, and that Amphan still caused large economic damage. Saving lives is a necessary first stage of resilience; protecting incomes and assets is a second, harder stage.

6.2 The Netherlands: planning for uncertain futures

The Netherlands has lived below sea level for centuries, but its modern approach to flood and water risk is distinctive for how it handles uncertainty about the future. The [Delta Programme](#) grew out of the recommendations of the Delta Commission and is anchored in a [Delta Act](#) that defines the powers of a Delta Programme Commissioner, sets out financing arrangements and obliges the government to update the programme every year. The programme has three parts: flood risk management, freshwater supply and spatial adaptation. All measures are financed from a dedicated Delta Fund, a separate item in the central government budget, which insulates long-term water safety from annual budget competition.

The method, known as Adaptive Delta Management, is the most developed example of planning under deep uncertainty in government. It works with scenarios rather than forecasts, develops adaptive strategies composed of flexible and robust measures, monitors conditions and recalibrates periodically. Scenarios are used to identify tipping points, the conditions at which present strategies stop working, which tells planners when larger works must begin. The design embodies the principles of pre-commitment (a legal mandate and a ring-fenced fund), learning (annual updating) and slack (explicit safety standards with margins built in). Its lesson for polycrisis is that institutional continuity matters as much as engineering: the Dutch approach survives changes of government because it is embedded in law and money.

6.3 Finland: comprehensive security as a whole-of-society habit

Finland's comprehensive security model treats resilience as a shared responsibility of authorities, businesses, non-governmental organisations and citizens. The [Security Committee](#) brings together permanent secretaries of the ministries, the Office of the President and the heads of key security agencies to develop a joint picture of risks and preparedness. The National Emergency Supply Agency works with the private sector through sector-specific pools, the insurance pool being one example, that share information, plan continuity and run exercises. Under the principle of locality, each actor retains its normal responsibilities in an emergency rather than handing them to a central authority. The model also reaches households through the concept of 72-hour preparedness, and it relies on regular exercises and defence courses that bring senior leaders from different sectors together.

Finland's case illustrates the enabling conditions in Figure 8: shared situational awareness, distributed authority, trust and learning loops. A caution is appropriate. Much of the available description of the model comes from the institutions themselves and from presentations by officials, and there is less independent measurement of how well it performs against real compound shocks. Its value for this report is as a mature example of institutional design built on the principles of polycentric governance.

6.4 Jamaica: pre-arranged finance meets a hurricane

In late October 2025, Hurricane Melissa struck Jamaica. The Government of Jamaica estimated total damages at between US\$6 billion and US\$9 billion, and the prime minister described preliminary damages and losses of US\$6–7 billion, close to a third of annual GDP. The World Bank [announced](#) that the storm triggered the full US\$150 million payout of the catastrophe bond it had arranged for the country, because the storm met the bond's pre-agreed parametric triggers. The payment was rapid because it did not require a damage assessment. [Reinsurance News](#) reports that the regional Caribbean Catastrophe Risk Insurance Facility, CCRIF, also indicated a record payout on Jamaica's tropical cyclone policy of about US\$70.8 million.

The example illustrates both the value and the limit of pre-arranged finance. The money arrived fast and without negotiation, which is exactly the property that matters in the first weeks of a disaster. But US\$150 million, together with the CCRIF payment, is a small fraction of the estimated damage. Pre-arranged instruments cover the liquidity gap and buy time; they do not replace the investment in prevention and the wider financing needed for reconstruction.

6.5 What the cases have in common

Table 6. Four cases compared

	Bangladesh	Netherlands	Finland	Jamaica
Core hazard	Tropical cyclones and storm surge	Flooding and freshwater shortage	Wide spectrum of security and supply crises	Hurricanes
Central mechanism	Volunteer warning network and shelters	Delta Act, Delta Fund, adaptive planning	Security Committee, supply pools, exercises	Parametric cat bond and regional insurance
Principles at work	Pre-commit, distribute authority, learn	Pre-commit, keep slack, learn	Distribute authority, build trust, learn	Pre-commit
Evidence of results	Deaths down by orders of magnitude	Long record of flood safety; annual programme updates	Strong institutional design; limited outcome data	Rapid payout after Melissa
Main limit	Livelihood and asset losses remain large	Designed for water, not all hazards	Outcomes less independently measured	Payout small relative to damage

H Heuristics synthesis of sources cited in section 6.

Four patterns stand out. First, each case invested before the shock, and the investment was sustained over decades rather than years. Second, each depended on institutional arrangements that outlast political cycles: a volunteer network, a law and a fund, a standing committee, a financing facility. Third, each distributed responsibility so that no single point of failure could stall the response. Fourth, each learned: the lessons of one disaster became design inputs for the next. None of these features is expensive by the standard of the losses it prevents, which is the subject of the next section.

7. The Economics of Resilience

Resilience has a reputation for being expensive and diffuse. The evidence says the opposite: well-targeted investment returns several dollars for each dollar spent, and the cheapest measures are among the least funded. The obstacle is not the arithmetic but the politics of paying now for losses that do not happen.

7.1 Returns on resilience investment

Two large bodies of work estimate the economics. The World Bank's [Lifelines](#) report (Hallegatte, Rentschler and Rozenberg, 2019) analysed resilient infrastructure in low- and middle-income countries across power, water, transport and telecommunications. It found that investing in more resilient infrastructure yields average net benefits of about US\$4.2 trillion, equivalent to about US\$4 of benefit for each US\$1 invested when climate change is taken into account. Without climate change the median benefit is closer to US\$2 per US\$1, and the authors stress that results depend on scenarios. A central message of the report is that resilience is cheaper and more effective when planners look at systems and users rather than at individual assets.

The Global Commission on Adaptation's [Adapt Now](#) report estimated that investing US\$1.8 trillion globally between 2020 and 2030 in five areas (early warning systems, climate-resilient infrastructure, improved dryland agriculture, mangrove protection and water resource resilience) could generate US\$7.1 trillion in total net benefits. It reported benefit-cost ratios ranging from 2:1 to 10:1 and, in some cases, higher, and described a triple dividend: avoided losses, economic gains from innovation, and social and environmental benefits. The [UNEP release](#) summarises the findings.




Figure 10. Reported benefit-cost ratios for resilience investment.

H Heuristics, using [World Bank, Lifelines](#) and the [Global Commission on Adaptation](#). The Commission point estimate is the ratio of US\$7.1 trillion net benefits to US\$1.8 trillion cost, computed by H Heuristics; the Commission itself reported a range of 2:1 to 10:1.

These ratios should be read with care. They are medians or ranges across scenarios, they depend on assumptions about hazards and discount rates, and the benefit categories differ between studies. The honest summary is not that every resilience project pays for itself, but that well-chosen investments typically do, with returns often well above the threshold at which governments normally approve infrastructure.

7.2 Early warning: the best-value place to start

Early warning stands out for its combination of low cost and high return. Coverage of the Global Commission's work reports that just 24 hours' notice of an impending hazard can cut ensuing damage by about 30 per cent, and that an investment of US\$800 million in early warning systems in developing countries could avoid losses of US\$3 billion to US\$16 billion per year. These figures are as reported in press coverage of the Commission's analysis and should be verified against its publications before formal citation.

Progress has been substantial and uneven. The UNDRR–WMO [Global Status of Multi-Hazard Early Warning Systems 2025](#), released at COP30, counts 119 countries, or 60 per cent of all countries, as having a multi-hazard early warning system, a 113 per cent increase over ten years. Coverage is highest in Asia-Pacific, at 72 per cent, and lowest in the Americas and Caribbean among the regions reported, at 51 per cent. Small island developing states have the lowest coverage among groups of countries in special situations, at 43 per cent, and Africa, though it has improved most since 2015, still has the lowest comprehensiveness scores. The Early Warnings for All initiative aims to protect everyone by 2027.




Figure 11. Multi-hazard early warning coverage by group, 2025.

H Heuristics, using the [UNDRR–WMO Global Status of Multi-Hazard Early Warning Systems 2025](#). Shares are of countries in each group reporting a multi-hazard system.

7.3 Paying for it: layers of risk finance

A resilient financing strategy matches instruments to the size and frequency of the loss, a practice known as risk layering. Table 7 sets out the layers; the Jamaican case in section 6 illustrated one of them.

Table 7. Layering risk finance: matching instruments to losses

Layer	Typical instruments	Speed	Best for	Limitation
Prevention and retention	Budget lines for prevention; reserve and contingency funds	Immediate	Frequent, small losses; risk reduction	Competes with other priorities each year
Contingent credit	Pre-approved credit lines; disaster drawdown options	Days to weeks	Medium losses; bridging liquidity	Adds to debt; terms matter
Risk transfer	Parametric insurance, risk pools, catastrophe bonds	Days	Severe, infrequent losses	Basis risk; premiums; weak on slow-onset events
Debt flexibility	Debt-pause clauses linked to disasters	Automatic	Preventing a disaster turning into a debt crisis	Needs lender buy-in across creditor groups

H Heuristics synthesis. Instrument availability varies by country; the Jamaica case in section 6.4 illustrates the risk-transfer layer.

Beyond these layers, post-disaster reconstruction finance and donor support arrive over months and are negotiated after the event, which is why the earlier layers matter. The protection gap shown in Figure 5

explains the urgency. When 80–90 per cent of catastrophe losses in emerging economies are uninsured, the burden falls on public budgets and households, and where public budgets are already stretched by debt service, as section 3 showed, it falls on whoever is least able to bear it. Swiss Re's estimate that insured losses could approach US\$400 billion by 2030 as exposure accumulates is a reminder that the demand for risk finance is rising, not falling.

7.4 Why societies underinvest anyway

If returns are high, why is spending low? Four reasons recur in the literature and in practice. First, the benefits of resilience are counterfactual: a flood that did not destroy a city leaves no photograph. Second, the time horizon of investment exceeds the political cycle, so the costs fall on one government and the benefits on another. Third, budgets are fragmented: the ministry that pays for a flood barrier is not the one that saves from reduced health costs. Fourth, the countries most exposed are those with the least fiscal room, so the choice is not between paying now and paying later but between paying now and not being able to pay at all.

These barriers point to remedies that are institutional rather than financial. The Dutch Delta Fund ring-fences water-safety spending from annual budget competition. Pre-arranged finance moves the decision from the crisis to the calm. Cross-sector appraisal captures benefits wherever they land. Section 9 returns to these institutional design choices.

8. A Policy Playbook for Six Groups of Actors

Resilience is not the job of a single ministry or a single level of government. This section sets out priorities for six groups whose actions together determine how a society handles compound shocks. The priorities follow from the principles in section 5 and the evidence in sections 3 to 7.

8.1 National governments

National governments set the legal framework, hold the largest balance sheet and own the relationships with international creditors. Their first task is to create a cross-risk capacity at the centre of government, whether a unit, a committee or a commissioner, with a mandate to look across sectors, publish a national risk assessment and commission regular stress tests of compound scenarios. Their second is to hold buffers deliberately: dedicated reserve funds, strategic stocks of critical goods, and fiscal rules that allow for counter-cyclical use of reserves. Their third is to arrange finance in advance: layered risk financing, disaster clauses in new debt contracts and standing arrangements with multilateral lenders. Their fourth is to protect the lifelines (energy, water, telecoms, payments, food distribution, health), with common standards that cover the interfaces between them.

8.2 Cities and subnational governments

Cities are where hazards, density and dependency meet, and where most practical decisions on land use, building standards and local services are taken. Priorities include mapping local critical nodes, building local early warning and evacuation capacity, keeping green and blue infrastructure that reduces flood and heat risk, and designing neighbourhood-scale backup for essential services. Cities are also natural laboratories for polycentric governance, because they can organise neighbourhood networks of the kind that made Bangladesh's volunteer system effective.

8.3 Firms and infrastructure operators

Operators of critical infrastructure and large firms control much of what actually fails in a cascade. Their priorities are to understand their own dependencies on suppliers and utilities beyond the first tier, to hold meaningful inventories or alternative sources for critical inputs, to design for graceful degradation rather than all-or-nothing operation, and to join sector-wide information-sharing and exercise arrangements of the kind Finland's supply pools provide. Boards should treat resilience as a core strategic issue and not as a cost to be minimised.

8.4 The financial sector

Banks, insurers, asset managers and pension funds are both exposed to cascades and capable of transmitting them. Their priorities are to include physical climate and compound-shock scenarios in stress tests, to build insurance products that expand cover in emerging economies, and to price risk transparently so that exposures are visible. Insurers' role in closing the protection gap, which Swiss Re puts at 80–90 per cent of losses in emerging economies, is among the most valuable contributions the sector can make. Supervisors should watch for concentration of exposures in common stresses.

8.5 Communities and civil society

The Bangladesh and Finland cases both show that communities are not just recipients of resilience policy but the carriers of it. Community networks supply local knowledge, trust and the last mile of communication that formal systems cannot reach. Priorities are to support volunteer and mutual-aid structures, to ensure that warnings reach and are understood by the most vulnerable, to practise drills, and to build trust in institutions, which research on polycentric governance identifies as the precondition for voluntary compliance in emergencies.

8.6 The international system

No country can solve compound shocks alone, because many of the common stresses and domino effects cross borders. The international priorities are to close the early-warning gaps identified by WMO and UNDRR, to improve the speed and predictability of crisis finance so that countries are not negotiating while losses mount, to make debt architecture more responsive to disasters, and to share technical standards for interdependent infrastructure. The Global Tipping Points Report's warning that tipping points are a new type of threat for which current international structures were not designed applies to this agenda as a whole.

Table 8. Priority actions by actor and time horizon

Actor	First 12 months	1–3 years	Measure of progress
National government	Create cross-risk unit; publish national risk assessment; map critical nodes	Reserve funds; layered risk finance; joint infrastructure standards	Months of buffer; share of risk pre-financed
Cities	Local dependency map; warning and evacuation plans	Backup for essential services; green-blue infrastructure	Share of population covered by local warning
Firms and operators	Tier-2 dependency audit; scenario exercise	Dual sourcing at critical nodes; graceful degradation	Recovery time after test disruptions
Financial sector	Compound-shock stress tests; protection-gap review	New insurance products; transparent risk pricing	Insured share of losses; stress-test results
Communities	Volunteer networks; drills; vulnerable-group outreach	Trusted local communication channels	Warning reach and comprehension
International system	Early-warning gap closure; faster crisis finance	Disaster-responsive debt terms; shared standards	Coverage of early warning; speed of finance

H Heuristics synthesis.

9. Governance, Measurement and a Phased Roadmap

A resilience strategy fails less often for lack of ideas than for lack of an institution that owns it, a way to know whether it is working and a sequence that makes it achievable. This section addresses each in turn.

9.1 Governance: who owns the cross-cutting problem

The core governance challenge is that polycrisis cuts across ministries, levels of government and the public-private divide. Three design choices recur in the cases reviewed. The first is a mandate with legal force. The Dutch Delta Act and the Finnish Security Committee both give a body a defined remit that persists beyond changes of government. The second is a dedicated, protected source of funds, as with the Delta Fund. The third is a requirement to report and update: the Delta Programme is revised every year and presented to Parliament, which creates a regular occasion for public scrutiny.

A central coordinating body should not centralise delivery. The evidence from polycentric governance is that resilience is stronger when many actors hold clear responsibilities and coordinate through shared awareness. The centre's job is to assemble the picture, set standards, test the system and resolve conflicts between sectors; the line ministries, cities, operators and communities do the rest.

9.2 Measurement: a short dashboard

What gets measured gets managed, and resilience has suffered from a lack of simple indicators. Table 9 proposes a short dashboard of twelve indicators in six groups. They are deliberately practical: each can be assembled by a government or large organisation from existing data within a year, and each moves in a direction that is clearly good or bad.

Table 9. A resilience dashboard: twelve indicators in six groups

Group	Indicator	What it shows	Good direction
Buffers	Months of essential import and service cover in reserves	Capacity to absorb a shock without emergency borrowing	Up
Buffers	Share of critical-service capacity held as spare or backup	Slack at critical nodes	Up
Coupling	Number of identified single points of failure	Hidden dependency risk	Down
Coupling	Share of critical systems with tested islanding or isolation	Real firebreaks	Up
Preparedness	Population covered by multi-hazard early warning	Warning reach	Up
Preparedness	Share of critical infrastructure meeting resilience standards	Lifeline protection	Up
Finance	Share of expected disaster losses covered by pre-arranged finance	Speed and certainty of funding	Up
Finance	Public interest payments as share of revenue	Fiscal room for response	Down
Learning	Share of national scenarios and drills run in the past year	Practice and testing	Up
Learning	Median time to recover critical services after disruptions	Recovery speed (also an early-warning signal)	Down

Group	Indicator	What it shows	Good direction
Equity	Coverage of warnings and cover for the poorest quintile	Fair distribution of protection	Up
Equity	Share of losses borne by households without insurance	Protection gap	Down

H Heuristics proposal. Indicators are illustrative; thresholds are not proposed because appropriate levels vary.

9.3 A phased roadmap

A realistic programme starts with understanding and low-cost foundations, builds capacity next, and then embeds and transforms. Figure 12 shows the phasing.




Figure 12. A three-phase roadmap across five workstreams.

H Heuristics. Conceptual roadmap; timings are indicative and will vary by country and sector.

Phase 1: Foundations (first 12 months). Build the system map and risk register, stand up a cross-risk unit with a clear mandate, audit critical nodes, close the most obvious early-warning gaps and establish a first tranche of reserves and insurance. These actions are cheap and inform every later decision.

Phase 2: Build (years 1 to 3). Run annual compound-shock stress tests, scale pre-arranged finance to cover the main modelled losses, begin systematic protection of lifelines and begin building firebreaks at the critical nodes the audit identified. Governance should mature into a body with legal force and a protected budget.

Phase 3: Embed and transform (years 3 to 10). Make resilience routine in planning, procurement and appraisal. Build modular grids and stockpiles, extend polycentric networks and, where adaptation is no longer enough, begin structural change: managed retreat, new energy, food and water systems, and institutional redesign.

10. Trade-offs, Risks and the Limits of This Analysis

A resilience agenda that does not state its costs and uncertainties is a slogan. This section sets out the main trade-offs, the ways the approach can go wrong and the evidence that would change the argument.

10.1 Trade-offs

Efficiency versus slack. Buffers and redundancy cost money in normal times. A grid with spare capacity, a firm with extra inventory and a government with idle reserves will look less efficient than its leaner counterparts until the day it matters. The case for slack is probabilistic and rests on tail events, which makes it politically fragile. The remedy is to make the cost explicit and the benefit measurable, for example through stress-test results.

Connectivity versus containment. Connections bring real benefits: trade raises incomes, grids share load, and finance spreads risk. The simulation in section 4 shows that connectivity can both protect and endanger. The aim is not to disconnect but to connect with care: to keep the benefits while installing firebreaks at the nodes where failure would be most damaging.

Speed versus legitimacy. Pre-commitment and delegated emergency powers speed up the response, but they concentrate authority. Resilience that is bought at the cost of public trust is brittle, since compliance in emergencies depends on legitimacy. Emergency powers should be tied to explicit triggers, sunset clauses and after-the-fact review.

Protection versus transformation. Money spent defending the present system is not available for building a different one. In some places, defence is the right choice; in others, such as locations that cannot be made safe, it delays an unavoidable change and raises its eventual cost. Planners need explicit criteria for when to switch from adaptation to transformation, and the Dutch approach of identifying tipping points in advance is a useful model.

Equity. Resilience investment tends to flow where assets and political voice are concentrated. Without deliberate attention, early warning, insurance and infrastructure protection will reach the better-off first, leaving the poorest, who are the most exposed, with the least cover. The dashboard in section 9 includes equity indicators for this reason.

10.2 Ways the agenda can fail

- **Resilience as an alibi.** A focus on adapting to shocks can be used to avoid addressing their causes. Resilience complements mitigation of emissions and other drivers; it does not substitute for it.
- **Over-centralisation.** A cross-risk unit that becomes a bottleneck or a source of command-and-control directives would undermine the distributed capacity the framework relies on.
- **Checklist compliance.** Dependency maps, risk registers and stress tests can become paperwork. The test of success is whether they change decisions and budgets.
- **False precision.** Dashboards and models give an impression of control. The scenarios in this report are illustrative, and real systems are less tidy.

- **Fighting the last crisis.** Resilience built against known hazards may leave a society exposed to a novel one. This is why the framework emphasises generic capacities (slack, firebreaks, learning) over hazard-specific fixes.

10.3 Limits of the evidence

Several limits should be stated plainly. The simulation in section 4 is a stylised model: it illustrates mechanisms but is not calibrated to any real system and should not be used to estimate real cascade risks. Death tolls and damage figures for historical disasters are contested, and this report uses ranges and approximations where sources differ. Some of the evidence for the benefits of early warning comes from press coverage of institutional analyses and should be checked against primary publications before formal citation. Case studies are chosen because they illustrate success and carry a selection bias, and the Finnish case in particular rests largely on descriptions from the institutions concerned. Finally, the report relies on published sources through October 2026, and some of the figures, particularly on debt, are drawn from secondary reporting of primary data that was not directly accessible.

10.4 What would change this argument

The argument would be weakened by evidence that: (i) losses in compound shocks are, on average, no larger than the sum of the separate shocks, which would undermine the entanglement thesis; (ii) the returns on resilience investment fall well below those reported when estimated from ex-post data rather than model scenarios; (iii) distributed, polycentric systems perform worse than centralised ones under compound stress; or (iv) pre-arranged finance fails to speed up recovery in practice. Each is testable, and a research programme that tracked outcomes across the cases in section 6 would help to settle them.

11. Conclusion: Building the Capacity to Keep Going

The central lesson of this report can be put in one sentence: the societies that fare best in an era of cascading shocks are not those that predict the next crisis, but those that have arranged themselves so that no single failure can bring them down.

The evidence assembled here shows an environment in which that arrangement matters more than it used to. Seven of nine planetary boundaries have been crossed, the first climate tipping point has been passed, disaster costs are about ten times what official figures record, and debt service is consuming the fiscal room on which crisis response depends. The experts surveyed by the World Economic Forum expect more turbulence, not less.

The evidence also shows that outcomes are not predetermined. Bangladesh cut cyclone deaths by orders of magnitude with volunteers, shelters and warnings. The Netherlands plans for an uncertain future with a law, a fund and a method for adapting. Finland has made resilience a shared habit. Jamaica received US\$150 million within days of a hurricane because it had arranged the money in advance. Each case combined early investment, distributed responsibility and learning.

For governments, the immediate agenda is modest and practical: map the dependencies, create a cross-risk capacity with a legal mandate, hold buffers on purpose, arrange finance before the shock, protect the lifelines and run the stress tests. For firms and communities, it is to understand dependencies beyond the first tier and to build trusted networks that work when formal systems do not. For the international system, it is to close the early-warning gaps and make crisis finance fast and predictable.

None of this is cheap, but none of it is as expensive as the alternative. The measure of a resilience strategy is not whether it prevents every crisis, which no strategy can, but whether it ensures that when shocks come, as they will, and when they arrive together, as they increasingly do, the society that absorbs them keeps its functions, its trust and its direction.

Sources and Further Reading

Global polycrisis and risk landscape

Lawrence, M., Homer-Dixon, T., Janzwood, S., Rockström, J., Renn, O. and Donges, J. F. (2024). [Global polycrisis: the causal mechanisms of crisis entanglement](#). *Global Sustainability*. Summary at the [Stockholm Resilience Centre](#).

Tooze, A. (2022). Welcome to the world of the polycrisis. *Financial Times*, 28 October 2022; see also [Chartbook](#).

World Economic Forum (2026). [Global Risks Report 2026](#) and [press release](#).

World Bank (2026). [Global Economic Prospects, June 2026](#).

Earth system and climate

Potsdam Institute for Climate Impact Research and partners (2025). [Planetary Health Check 2025](#).

Richardson, K. et al. (2023). [Earth beyond six of nine planetary boundaries](#). *Science Advances*.

Global Tipping Points Report 2025. [Release summary](#); programme site at [global-tipping-points.org](#); coverage in [Time](#).

Armstrong McKay, D. et al. (2022). Exceeding 1.5°C global warming could trigger multiple climate tipping points. *Science*. [University of Exeter summary](#).

World Meteorological Organization (2026). State of the Global Climate 2025. [Summary at the Royal Meteorological Society](#); [WMO provisional release](#).

Intergovernmental Panel on Climate Change (2023). [Sixth Assessment Report: Synthesis Report](#).

Disaster costs, debt and finance

UNDRR (2025). [Global Assessment Report on Disaster Risk Reduction 2025: Resilience Pays](#); [PreventionWeb summary](#).

Swiss Re Institute (2026). [sigma 1/2026: Natural catastrophes](#) and [press release](#).

UNCTAD (2025). [A World of Debt](#).

Hallegatte, S., Rentschler, J. and Rozenberg, J. (2019). [Lifelines: The Resilient Infrastructure Opportunity](#). World Bank.

Global Commission on Adaptation (2019). Adapt Now: A Global Call for Leadership on Climate Resilience. [WRI release](#); [UNEP release](#).

UNDRR and WMO (2025). [Global Status of Multi-Hazard Early Warning Systems 2025](#).

Networks, cascades and resilience theory

Buldyrev, S. V., Parshani, R., Paul, G., Stanley, H. E. and Havlin, S. (2010). [Catastrophic cascade of failures in interdependent networks](#). *Nature* 464.

Helbing, D. (2013). [Globally networked risks and how to respond](#). *Nature* 497.

Watts, D. J. (2002). [A simple model of global cascades on random networks](#). *Proceedings of the National Academy of Sciences* 99.

Haldane, A. G. and May, R. M. (2011). [Systemic risk in banking ecosystems](#). *Nature* 469.

Scheffer, M. et al. (2009). [Early-warning signals for critical transitions](#). *Nature* 461.

Holling, C. S. (1973). [Resilience and stability of ecological systems](#). *Annual Review of Ecology and Systematics* 4.

Walker, B., Holling, C. S., Carpenter, S. R. and Kinzig, A. (2004). [Resilience, adaptability and transformability in social-ecological systems](#). *Ecology and Society* 9(2).

Folke, C. et al. (2010). [Resilience thinking: integrating resilience, adaptability and transformability](#). *Ecology and Society* 15(4). See also the [Stockholm Resilience Centre](#).

Ostrom, E. (2010). [Beyond markets and states: polycentric governance of complex economic systems](#). *American Economic Review* 100(3).

Case studies

ENTSO-E expert panel (2026). Final report on the 28 April 2025 Iberian blackout. Summaries: [pv magazine](#); [T&D World](#).

FERC and NERC (2021). Final report on the February 2021 freeze. [T&D World summary](#); [Power Engineering on FERC standards](#).

[Dialogue Earth](#) and [The Business Standard](#) on Bangladesh's cyclone preparedness.

Government of the Netherlands. [Delta Programme](#) and the [Delta Act](#).

Finnish Security Committee. [Concept of Comprehensive Security: Building National Resilience in Finland](#).

World Bank (2025). [Hurricane Melissa triggers 100% payout of US\\$150 million catastrophe bond for Jamaica; Reinsurance News](#).

Appendix: Methods, Data and the Illustrative Simulation

A.1 Approach and status of the figures

This report synthesises published institutional and peer-reviewed research. Every quantitative claim about the world is attributed in the text to its source, and the links in the text and in the Sources section point to the documents or summaries used. Where only a secondary summary of a primary document could be accessed, the text says so. Table A1 records the status of each figure.

Table A1. Status of each figure

Figure	Type	Data basis
1 Planetary boundaries	Reported data (binary)	Planetary Health Check 2025
2 Three pathways	Conceptual schematic	Lawrence et al. (2024); illustrative examples
3 Coupling web	Conceptual schematic	Qualitative judgement; not measured
4 Warming and thresholds	Reported data	WMO; Global Tipping Points Report 2025; IPCC AR6
5 Disaster costs and insurance	Reported data	UNDRR GAR 2025; Swiss Re sigma 1/2026
6 Debt buffers	Reported data	UNCTAD, as reported; 2025 figure rounded
7 Cascade simulation	Illustrative simulation	H Heuristics model; see A.2
8 Framework	Conceptual schematic	Resilience literature
9 Bangladesh	Reported data (approximate)	Press and institutional summaries; tolls contested
10 Benefit-cost ratios	Reported data	World Bank; Global Commission on Adaptation
11 Early warning coverage	Reported data	UNDRR–WMO 2025
12 Roadmap	Conceptual schematic	H Heuristics

H Heuristics.

A.2 The cascade simulation

Figure 7 is generated by a threshold-cascade model of the type introduced by Watts (2002). The model has $N = 1,000$ systems connected in an undirected network. Each system has a tolerance ϕ , and it fails when the share of its connected neighbours that have failed is at least ϕ . Cascades start from one or more systems selected at random to fail, and failure spreads in rounds until no further system fails. The outcome measure is the share of all systems that have failed at the end of the cascade, averaged over repeated runs.

In panels A and B, networks are random (Erdős–Rényi) graphs with a given average number of links per system, ranging from 1 to 12; each point averages 60 networks with new random initial failures. Panel A uses tolerances of 0.12, 0.18 and 0.30 with one initial failure. Panel B uses a tolerance of 0.18 with one or ten simultaneous initial failures. In panel C, the 1,000 systems are divided into 20 modules of 50; the average number of links per system is 4.5; a given share of all links is placed between modules and the remainder within modules; the tolerance is 0.18; each point averages 80 networks with one initial failure. The random-number generator is seeded for reproducibility.

The model deliberately omits many features of real systems: heterogeneous load and capacity, directed and weighted links, repair, behavioural responses and external shocks that vary in size. Its purpose is to illustrate the qualitative effects of tolerance, connectivity, simultaneous shocks and modular structure. It should not

be used to estimate the failure risk of any specific real system. The danger zone in connectivity depends on tolerance and on network structure, and in real, more heterogeneous systems it may shift or disappear.

A.3 Reporting conventions

Dollar figures are US dollars. “Per cent” and “billion” follow British and international usage, where a billion is a thousand million and a trillion is a thousand billion. Where sources give ranges for historical death tolls or losses, the report shows approximations and says so. All web links were tested in October 2026. Several publishers (including UNDRR, UNCTAD, the World Economic Forum, Swiss Re and Dialogue Earth) block automated access, so those pages were confirmed through search results rather than opened directly; where a primary document could not be read, the link goes to the most authoritative accessible summary.