

L-ARC™ FRAMEWORK: EMPIRICAL CROSS-HAZARD VALIDATION

Nepal 2026 Case Study and Comparative Evidence Across Six Hazards

Research White Paper | Testing the Intelligence -> Governance -> Capital resilience
mechanism across six real-world hazards

Central finding

Evidence across six materially different hazards is consistent with a repeatable L-ARC™ resilience mechanism: risk intelligence becomes protective when it reaches legitimate decision-makers in time, authorised action is backed by deployable capacity, and resilience investment strengthens the next cycle of intelligence and preparedness. Nepal 2026 provides the strongest within-event trace in this portfolio, while five comparator cases test whether the same institutional logic remains analytically useful across different hazards.

Dr. Surajit Lahiri

Creator, Lahiri AI-Resilience Continuum (L-ARC™)

Evidence status

This report presents an evidence-based retrospective test of the L-ARC™ mechanism across flood, wildfire, extreme heat, pandemic, tropical cyclone and critical-infrastructure failure. The same core architecture is applied across all six cases. The present numerical scores are author-generated, illustrative research outputs intended to test framework coherence rather than establish definitive institutional rankings. Universal calibration, comparative superiority and prospective causal effect sizes remain questions for independent testing.

Evidence cut-off: Nepal casualty and loss figures use the cited 4 September 2026 official situation update

EXECUTIVE SUMMARY

White Paper at a Glance

This white paper asks whether the same L-ARC™ institutional mechanism remains analytically useful when applied, without changing its core architecture, across materially different crises. It is a retrospective, evidence-led stress test of framework coherence and diagnostic utility - not a claim of universal causal efficacy.

Research question	Does the Intelligence -> Governance -> Capital mechanism remain meaningful across different hazards, timescales and governance settings?
Evidence base	Nepal 2026 as the anchor case, with Maui wildfire 2023, European heat 2022-23, India COVID-19 second wave 2021, Super Typhoon Rai/Odette 2021, and the Texas grid failure 2021.
Principal result	Risk knowledge becomes protective when it reaches legitimate decision-makers in time, authorised action can mobilise capacity, and resilience investment strengthens future intelligence and preparedness.
Strongest empirical signal	Interface performance: the hand-offs from intelligence to governance, governance to resources, and resilience investment back to stronger future intelligence.
Claim boundary	The study supports retrospective mechanism coherence and diagnostic applicability. It does not establish a universal effect size, exact losses avoided, universal scoring thresholds, predictive superiority, or institutional endorsement.

Why this matters

Across the six cases, institutions often had some risk knowledge, formal authority or response capability. The recurring weakness was conversion: whether information remained available, reached a legitimate decision pathway with usable lead time, mobilised resources, and was reinforced through durable investment. L-ARC™ makes these institutional hand-offs explicit rather than treating resilience as a single undifferentiated capability.

The L-ARC™ institutional lens

Anticipatory Intelligence	Do we know enough, early enough and reliably enough to support preventive action?
Governance Legitimacy	Can recognised actors legitimately, accountably and defensibly act before certainty is complete?
Capital Sustainability	Can resilience capability be financed, maintained and strengthened over time?
Interfaces	Does intelligence actually become action, does action mobilise capacity, and does investment strengthen the next cycle?

Nepal 2026: four signals that make the mechanism visible

~14 min	>916	~41 min	>=4
School warning lead time [4]	Students and staff evacuated before destruction [4]	Border-impact to public-alert interval [5]	Gauges reported lost in the wider system [5]

Executive Summary

This report asks a demanding institutional question: does the same L-ARC™ mechanism remain visible when the framework is applied, without changing its architecture, to materially different crises? The cases are the 26 August 2026 Bhote Koshi-Trishuli disaster in Nepal; the 2023 Lahaina wildfire in Maui; the 2022-2023 European heat crisis; India's 2021 COVID-19 second wave; Super Typhoon Rai/Odette in the Philippines in 2021; and the Texas electricity-grid failure of 2021.

The evidence is consistent with the L-ARC™ mechanism within the boundary of a retrospective case-study design. Across all six cases, the same institutional problem recurs: hazard knowledge does not automatically become resilience. The decisive variable is conversion - whether credible intelligence reaches a legitimate decision pathway in time, whether authorised action can mobilise resources, and whether resilience investment strengthens future sensing, preparedness and institutional memory.

What the evidence supports

The six-case analysis provides evidence consistent with cross-hazard retrospective mechanism coherence and cross-geographic diagnostic applicability for L-ARC™. Nepal provides a direct warning-decision-action trace; the Philippines supplies a positive large-scale forecast-to-evacuation contrast; Maui and Texas show different conversion and reinforcement failures. The evidence does not establish a universal causal effect size, universal scoring thresholds, predictive superiority or institutional endorsement.

The report is structured for institutional use. Nepal is the primary case study because it contains a documented success-and-failure contrast within the same event. Five comparator hazards test whether the L-ARC™ mechanism remains coherent across different physical processes, timescales and governance settings. A common measurement architecture is used throughout, with observed facts, derived arithmetic, interpretive scores and counterfactual claims kept explicitly separate.

Nepal is the anchor case

Nepal provides the strongest within-event test. The Government of Nepal's 4 September situation update recorded 1,287 bodies recovered, around 5,083 people still missing and 13,093 people rescued, including more than 300 foreign nationals [1]. NDRRMA's initial multi-sector assessment placed total losses at NPR 387.54 billion (about US\$2.6 billion) and immediate early-recovery and initial rehabilitation needs at NPR 7.95 billion (about US\$53 million) for the first four months [1]. UNDP's RAPIDA assessment estimated 84,270 affected people across 17 municipalities in six districts, with 76% living in municipalities classified as having above-average pre-flood vulnerability; preliminary potential building damage alone was US\$158.6 million ±30% [2]. UNDP also reported 2.2 million tonnes of broader debris and 901.1 MW of operating or planned energy capacity affected when reported operating and construction categories are combined [3].

The most probative mechanism evidence is local. At Tribhuvan Trishuli Secondary School, an upstream telephone warning provided roughly 14 minutes of lead time. The headmaster stopped classes, redirected buses and evacuated more than 900 students and 16 staff before the school was destroyed [4]. In the wider system, Reuters reported that at least four gauges were lost and that a public alert followed the border impact by about 41 minutes [5]. The contrast does not prove how many national deaths were preventable. It does demonstrate that minutes, decision rights and warning continuity were operational variables in the same event.

Cross-hazard result

Case	All	GLI	CSI	GPI	LCCS*	Maturity	Evidence confidence
Nepal flood 2026	42	58	42	58	46.8	Level 3	Moderate-High
Maui wildfire 2023	46	34	34	33	37.6	Level 3	High
Europe heat 2022-23	82	58	54	60	63.6	Level 3	Moderate
India COVID-19 2021	50	54	38	55	46.8	Level 3	Moderate-High
Philippines Rai 2021	74	70	54	70	65.4	Level 4	High
Texas grid 2021	50	38	38	35	41.6	Level 3	High

* LCCS is the provisional geometric-mean research metric from the L-ARC™ Master Specification. In this paper it is secondary to the underlying capability profile and must not be interpreted as disaster probability, loss probability or certification.

The pattern is not that every crisis has low intelligence. Europe has the strongest Anticipatory Intelligence profile in the portfolio, yet high mortality persisted. The Philippines provides the clearest positive contrast: strong forecast-to-evacuation conversion moved more than 400,000 people from the storm's path before landfall [24]. Texas provides the clearest learning failure: a similar 2011 inquiry produced recommendations that FERC later said were not acted upon before the 2021 freeze [27].

ABSTRACT AND CLAIM BOUNDARY

Abstract

The Lahiri AI-Resilience Continuum (L-ARC™) defines resilience as continuity across three co-dependent institutional capabilities: Anticipatory Intelligence, Governance Legitimacy and Capital Sustainability. This institutional case-study report subjects that architecture to a cross-hazard retrospective stress test using six cases spanning flood, wildfire, extreme heat, pandemic, tropical cyclone and critical-infrastructure failure. The same measurement architecture is applied to each case: All, GLI, CSI, GPI, the IGC/GCR/CIR interface metrics, and the six-level maturity model. Evidence is drawn principally from government, United Nations, intergovernmental, regulatory and official investigation sources, with Reuters used where it provides uniquely documented event timelines or interviews.

The analysis finds evidence consistent with the L-ARC™ architecture across all six hazard classes. The most recurrent pattern is an interface effect: meaningful risk information frequently existed before the worst consequences, yet resilience performance depended on whether that information became timely and defensible action, whether action was backed by deployable capacity, and whether long-term investment had hardened the systems required for the next event. Nepal supplies a strong within-event contrast between a successful local Intelligence-to-Governance hand-off and broader monitoring/warning discontinuity. The Philippines shows the positive direction of the mechanism, while Texas shows failure of institutional learning and capital reinforcement. Together, the cases support retrospective mechanism coherence and diagnostic utility while leaving universal calibration and prospective causal effect sizes open for broader testing.

Claim boundary

- Supported in this study: repeatable retrospective mechanism fit, cross-hazard diagnostic applicability, interface relevance, and positive/negative case discrimination.
- Strongly supported: decision latency, continuity of intelligence, governance permission and durable capital as measurable resilience variables.
- Not established: a universal L-ARC effect size, universal scoring thresholds, exact deaths or losses avoided, or superiority over Sendai, EW4All, AWARE or other established frameworks.
- Not claimed: that L-ARC™ predicted any of the six physical hazards or that AI can deterministically predict rare fast-onset events.

- Further scientific strengthening: independent raters and inter-rater reliability testing, matched-framework comparator coding, larger cross-case replication, construct-validity testing and prospective pilots.

Institutional summary

Cross-hazard retrospective analysis provides evidence consistent with the L-ARC™ core mechanism: anticipatory information becomes protective when it crosses into legitimate, timely institutional action; durable resilience further depends on whether authorised action can mobilise capacity and whether investment reinforces future intelligence.

Paper structure

1. Research question and design
 2. L-ARC™ methodology used without modification
 3. Nepal 2026 - primary real-world stress test
 4. Maui 2023 - wildfire
 5. Europe 2022-2023 - extreme heat
 6. India 2021 - COVID-19 second wave
 7. Philippines 2021 - Super Typhoon Rai/Odette
 8. Texas 2021 - grid failure
 9. Cross-case results
 10. Test of the six L-ARC™ propositions
 11. International alignment and framework differentiation
 12. Robustness, falsifiability and claim boundaries
 13. Institutional adoption, comparator testing and independent replication
 14. Conclusion
- References and evidence appendices

RESEARCH DESIGN AND METHOD

1. Research question and design

The principal research question is: Does the L-ARC™ architecture provide a coherent and useful explanation of resilience performance across materially different real-world hazards when the methodology is held constant?

Three subsidiary questions are tested: (1) do the same three foundational capabilities remain meaningful across different hazards; (2) do failures repeatedly occur at the interfaces between those capabilities; and (3) can the framework distinguish cases where warnings were successfully converted into preventive action from cases where intelligence existed but conversion failed?

1.1 Case-selection logic

The portfolio deliberately includes different physical mechanisms, timescales and governance settings. It also includes a comparatively successful anticipatory-action case - the Philippines - to reduce the risk that L-ARC™ simply assigns low scores after failed outcomes. Nepal is treated as the primary case because the event is recent, richly documented, transboundary and contains a nested success/failure contrast within the same hazard.

Hazard	Case	Unit of assessment
Flood / cascading mountain hazard	Nepal Bhote Koshi-Trishuli, 2026	Transboundary river-corridor and Nepal multi-agency warning/response system
Wildfire	Lahaina, Maui, 2023	Maui County/state wildfire preparedness, warning and response system
Extreme heat	Europe, 2022-2023	Regional heat-risk and health-preparedness ecosystem
Pandemic	India COVID-19 second wave, 2021	National/state surge and medical-oxygen preparedness system
Tropical cyclone	Philippines Typhoon Rai/Odette, 2021	National/local typhoon warning, evacuation and DRRM ecosystem
Infrastructure failure	Texas electricity grid, 2021	ERCOT/Texas electricity-natural-gas resilience system

Scores from these units are not treated as national league tables. They are structured retrospective profiles. Different assessment boundaries, evidence density and institutional scales limit direct comparability.

1.2 Evidence hierarchy

- Tier 1: government, regulator, United Nations, WMO/WHO, official investigation and official disaster-management records.
- Tier 2: peer-reviewed research and institutional technical assessments.
- Tier 3: Reuters reporting where it provides uniquely documented timelines, interviews or post-event policy information not yet available in primary records.
- Derived arithmetic is shown explicitly and is never presented as a reported source value.
- Interpretive L-ARC™ scores are separated from observed facts.

1.3 Positionality and conflict of interest

The author is the originator of L-ARC™ and therefore has an intellectual and reputational interest in the framework. The current case scores are also assigned by the framework creator. This conflict and methodological limitation are explicit. Safeguards used here are a fixed data cut-off, transparent scoring rules, inclusion of strong comparator frameworks, explicit negative findings, no fitted weights, and clear separation between reported facts, derived arithmetic, interpretation and untested counterfactuals. Independent replication and reliability testing are therefore treated as required next tests rather than optional refinements.

L-ARC METHODOLOGY

2. L-ARC™ methodology used without modification

The foundational continuum is Anticipatory Intelligence -> Governance Legitimacy -> Capital Sustainability -> stronger future intelligence. The operating model is Early Signals -> Signal Trust and Context -> AI Risk Intelligence -> Institutional Permission -> Capital Confidence -> Preventive Action and Durability -> learning. The cross-hazard test does not add a fourth pillar or alter the framework to fit individual cases.

2.1 Measurement architecture

Index	Dimensions	Core question
All - Anticipatory Intelligence Index	A1 coverage; A2 reliability; A3 anticipatory performance; A4 contextual intelligence; A5 explainability/actionability	Does the institution know enough, early enough and reliably enough to support preventive action?
GLI - Governance Legitimacy Index	G1 authority; G2 timeliness; G3 accountability; G4 explainability/trust; G5 legal, policy and ethical defensibility	Can the institution legitimately and accountably act on what it knows?
CSI - Capital Sustainability Index	C1 resilience investment; C2 preparedness/capability; C3 monitoring sustainability; C4 financial continuity; C5 resilience value/durability	Can resilience be financed, maintained and strengthened over time?
GPI - Governance Permission diagnostic	Authority, escalation, thresholds, risk tolerance, decision confidence, explainability, accountability, defensibility	Is the institution permitted and sufficiently confident to act before certainty is complete?

2.2 Retrospective evidence-coding rule

The master 0-100 scale is implemented as five ordinal evidence bands with midpoint coding. This avoids false precision while permitting transparent cross-case arithmetic.

Evidence band	Coding value	Interpretation
0-19	10	Capability absent or unsupported by credible evidence
20-39	30	Informal, fragmented or inconsistently applied
40-59	50	Partially defined and operational in limited settings
60-79	70	Formally established, evidenced and generally applied
80-100	90	Institutionalised, tested, reviewed and continuously improved

$$\begin{aligned}
 \text{All} &= (A1 + A2 + A3 + A4 + A5) / 5 \\
 \text{GLI} &= (G1 + G2 + G3 + G4 + G5) / 5 \\
 \text{CSI} &= (C1 + C2 + C3 + C4 + C5) / 5 \\
 \text{LCCS} &= (\text{All} \times \text{GLI} \times \text{CSI})^{(1/3)} \text{ [provisional research metric only]}
 \end{aligned}$$

Scores in this report are illustrative research outputs intended to test framework coherence rather than establish definitive institutional rankings. The midpoint values are coding anchors rather than universal thresholds. Because the current ratings are author-generated, they should be treated as transparent author-coded estimates until independent inter-rater reliability is established.

GPI scoring rule

GPI is a separate governance-permission diagnostic and is not included in LCCS. It is scored across eight equal-weight permission dimensions that correspond to the conceptual elements already defined in Section 2.1: formal authority (P1), escalation clarity (P2), trigger/threshold clarity (P3), risk tolerance for proportionate action under uncertainty (P4), decision confidence/evidentiary sufficiency (P5), explainability (P6), accountability/reviewability (P7), and legal-policy-ethical defensibility (P8).

Component	Operational coding question
P1 - Formal authority	Is a recognised actor formally or operationally empowered to decide and act?
P2 - Escalation clarity	Is the pathway for escalation, hand-off and decision ownership clear enough to operate within the available lead time?
P3 - Trigger/threshold clarity	Are decision triggers, thresholds or conditions for preventive action defined and usable?
P4 - Risk tolerance under uncertainty	Can proportionate preventive action be authorised before certainty is complete?
P5 - Decision confidence	Is the available evidence sufficiently trusted and decision-relevant to support action?
P6 - Explainability	Can the basis for the decision be explained to affected actors, reviewers or the public where required?
P7 - Accountability/reviewability	Are the decision, responsible actor and rationale attributable and open to appropriate review?
P8 - Defensibility	Is the contemplated action legally, policy-wise and ethically defensible?

GPI evidence anchors

Anchor	Interpretation
0	No credible evidence of usable permission, or evidence that action was effectively unavailable.
20	Nominal or highly fragmented permission; authority, escalation or defensibility is materially unclear.
40	Partial permission exists, but material ambiguity, delay, confidence or accountability constraints remain.
60	Permission is defined and usable in relevant settings, but execution remains conditional or inconsistent.
80	Permission pathways are clear, generally executable before certainty is complete, and supported by reviewable governance.
100	Permission is institutionalised, tested, timely, consistently executable and demonstrably defensible across relevant operating conditions.

$GPI = \text{round}[(P1 + P2 + P3 + P4 + P5 + P6 + P7 + P8) / 8]$. Each component is assigned one of the six anchors above. The arithmetic mean is rounded to the nearest whole number for reporting. Because GPI is a diagnostic rather than a component of LCCS, its anchor structure is kept separate from the midpoint-only AII/GLI/CSI coding rule. The component evidence trail should remain available so an independent coder can reproduce or challenge the reported GPI.

2.3 Interface metrics

IGC = eligible credible alerts receiving documented governance disposition / eligible credible alerts requiring governance consideration x 100

GCR = approved actions adequately resourced within required time / approved actions requiring resource mobilisation x 100

CIR = approved intelligence-strengthening investments made operational / approved intelligence-strengthening investments x 100

Formal percentages are calculated only when the historical record contains a defensible denominator. Where it does not, the paper reports an interface diagnosis - Strong, Moderate or Weak - and states that the formal ratio is not estimable. This avoids manufacturing precision from incomplete historical data.

2.4 Maturity classification decision rule

Maturity is not assigned from LCCS cut-points. It is a non-compensatory, evidence-based classification of end-to-end institutional behaviour. A high Anticipatory Intelligence score cannot by itself move a system to a higher maturity level when warning-to-action conversion is inconsistent, and a relatively modest composite score does not preclude Level 4 when systematic preventive conversion is directly evidenced. LCCS and maturity therefore answer different questions: LCCS summarises capability balance; maturity describes how reliably the continuum operates in practice.

Level	Name	Minimum decision criterion	Interpretive guardrail
1	Fragile	Core resilience capabilities are largely absent, unsupported or highly fragmented.	No evidence of a dependable end-to-end preventive pathway.
2	Reactive	Some capability exists, but institutional action occurs mainly after disruption or after preventive lead time has been lost.	Monitoring or plans alone do not qualify as anticipatory operation.
3	Aware	Risk is monitored and institutionally recognised, but at least one critical Intelligence-to-Governance, Governance-to-Capital or reinforcement interface remains inconsistent or incomplete.	High AI or LCCS cannot compensate for incomplete conversion.
4	Anticipatory	Credible early intelligence is repeatedly or systemically converted into timely, legitimate preventive action before impact, with resources mobilised to implement that action.	Requires observed preventive conversion, not simply good forecasts or formal plans.
5	Adaptive	Level 4 conditions plus documented learning that changes protocols, thresholds, investment or monitoring and is subsequently made operational.	Post-event recommendations must be implemented, not merely announced.
6	Resilient	Adaptive integration is sustained across cycles, with redundancy, continuity, accountability and continuous improvement institutionalised.	Requires longitudinal evidence across multiple events or operating cycles.

Application to this portfolio: the Philippines is classified at Level 4 because official warnings, pre-positioning and pre-emptive evacuation show system-level conversion before landfall, including more than 400,000 people evacuated [23,24]. Europe remains Level 3 at the regional aggregate despite a higher LCCS because strong forecasting coexisted with incomplete heat-health action-plan and surveillance coverage across jurisdictions [17]. The difference is therefore driven by observed conversion behaviour, not by a higher or lower composite score.

PRIMARY CASE: NEPAL 2026

3. Nepal 2026 - Bhote Koshi-Trishuli disaster

3.1 Facts before interpretation

On 26 August 2026, a catastrophic high-mountain failure generated a destructive flood and debris flow through the Bhote Koshi-Trishuli system. ESA satellite analysis reported a rapid slope failure associated with a glacier collapse on Langtang Lirung; the resulting torrent of water, ice and debris travelled nearly 100 km downstream [7]. The Government of Nepal continued to study the event's causes and evolution, and this paper therefore does not treat every aspect of the physical trigger as finally settled [1].

UNDRR characterised the disaster as a stark reminder of the transboundary nature of disaster risk, the fragility of mountain ecosystems and the difficulty of early warning and early action for cascading hazards in mountainous regions [8].

3.2 Quantified evidence ledger

Indicator	Observed / derived value	Source / calculation	Constraint
Bodies recovered	1,287	Government of Nepal, 4 Sep situation update [1]	Official snapshot at 14:00 NST; verification ongoing
People missing	around 5,083	Government of Nepal [1]	Official snapshot; evolving
People rescued	13,093	Government of Nepal [1]	Includes >300 foreign nationals
Affected population	84,270; 17 municipalities; 6 districts	UNDP RAPIDA [2]	Government-approved estimate
Above-average vulnerability	76% of affected-municipality population	UNDP MVI overlay [2]	Municipality classification, not individual status
Derived population in those municipalities	about 64,045	$84,270 \times 0.76$	Derived arithmetic
Potential building damage	US\$158.6m $\pm$ 30%	UNDP preliminary model [2]	Not total disaster loss
Derived uncertainty range	US\$111.0m-US\$206.2m	$158.6 \times 0.70 / 1.30$	Derived arithmetic
Broader debris estimate	2.2 million tonnes	UNDP [3]	Includes property, vegetation, rock, sediment
Operating generation stopped	431.1 MW	UNDP [3]	11 hydro + 1 solar plant
Projects under construction damaged	470 MW	UNDP [3]	15 projects
Derived energy capacity exposure	901.1 MW	$431.1 + 470$	Capacity exposure, not monetary loss
Initial total loss assessment	NPR 387.54bn (~US\$2.6bn)	Government of Nepal / NDRRMA [1]	Initial multi-sector assessment; comprehensive assessment pending
Immediate early-recovery need	NPR 7.95bn (~US\$53m), first 4 months	Government of Nepal / NDRRMA [1]	Early recovery / initial rehabilitation; not total reconstruction cost
School protective-action population	>916 people	>900 students + 16 staff [4]	Directly observed interview evidence
School lead time	about 14 minutes	Reuters [4]	Interview-based timing

Indicator	Observed / derived value	Source / calculation	Constraint
Public-alert interval after border impact	about 41 minutes	08:32 impact to 09:13 alert [5]	Different locations/processes; not causal equivalence

Why the 41/14 comparison matters - and why it must not be abused

A roughly 41-minute public-alert interval is about 2.93 times the 14-minute lead time that proved operationally useful at the school. The locations and processes are different, so this is not a causal equivalence. It does show that minutes were operationally meaningful in this event and that decision latency cannot be treated as an abstract variable.

3.3 Real-world test: Anticipatory Intelligence

Nepal did not lack monitoring. The central issue was continuity. Reuters reported that sensors and gauging stations existed, but monitoring was not extensive; at least four gauges were knocked out and officials did not immediately know that they had failed [5]. This is precisely the L-ARC™ distinction between sensor presence and anticipatory intelligence. A system must know what it is seeing, detect when it has gone blind, fuse signals with exposure and context, and communicate uncertainty in time for action.

All dimension	Score	Evidence interpretation
A1 Signal coverage	50	River monitoring and regional hazard intelligence existed; high-mountain source-hazard coverage remained incomplete.
A2 Signal reliability	30	Critical gauges were lost during the surge and failure awareness was incomplete [5].
A3 Anticipatory performance	30	For the fastest upstream corridor, system-level usable lead time was extremely limited.
A4 Contextual intelligence	50	Transboundary, glacier and cascading risk were recognised, but source-to-corridor integration remained incomplete.
A5 Explainability/actionability	50	Warnings were actionable once received, but uneven lead time constrained use across the corridor.

Nepal AI = 42/100 - Partially established

The finding is not that Nepal lacked intelligence. It had hydrometeorological knowledge and warning capability. The bottleneck was continuity and source-sensitive anticipation for an exceptionally rapid transboundary cascading hazard.

3.4 Real-world test: Governance Legitimacy and GPI

The school micro-case is the cleanest mechanism evidence because the critical variables are visible. A warning arrived; a person with recognised authority judged it credible enough to act; he did not wait for perfect confirmation; buses were redirected and people moved to safety; the school was then destroyed [4].

Signal	Legitimate decision	Protective outcome
Upstream telephone warning; roughly 14 minutes lead time	Headmaster stops classes, rings alarm, redirects buses and triggers evacuation	>900 students and 16 staff moved before school destruction

At this tightly bounded school-level unit, one eligible credible warning produced one documented governance action. Therefore school-level IGC = $1/1 \times 100 = 100\%$ (n=1). This is a mechanism illustration only and must not be extrapolated to Nepal's national warning system.

GLI dimension	Score	Evidence interpretation
G1 Authority clarity	70	National, district and technical response roles were identifiable; local school authority was clear.
G2 Decision timeliness	50	Once information arrived, actions could be rapid; system-wide information arrival was late in the most exposed corridor.

GLI dimension	Score	Evidence interpretation
G3 Accountability	70	Warning and response timelines are reconstructable; government coordination was visible.
G4 Explainability/trust	50	Warnings were actionable in some locations, but uneven lead time constrained practical trust and response.
G5 Legal/policy/ethical defensibility	50	Formal disaster governance existed; transboundary cascading-hazard governance and information continuity remained incomplete.

Nepal GLI = 58/100 | GPI = 58/100

Governance authority existed. The central constraint was not simply permission; it was whether actionable intelligence reached legitimate actors before residual lead time disappeared. Governance Legitimacy is therefore time-dependent.

3.5 Real-world test: Capital Sustainability

The Government of Nepal's 4 September situation update reported NDRRMA's initial multi-sector loss assessment at NPR 387.54 billion (about US\$2.6 billion), alongside NPR 7.95 billion (about US\$53 million) in immediate early-recovery and initial rehabilitation needs for the first four months [1]. These national figures are broader than UNDP's building-only model, which estimated potential building damage of US\$158.6 million $\pm 30\%$, an implied range of approximately US\$111.0 million to US\$206.2 million [2]. Energy exposure was also substantial: 431.1 MW of operating generation stopped and 470 MW of projects under construction were damaged, producing a combined reported operating/planned capacity exposure of 901.1 MW [3]. Reuters separately reported preliminary estimates of around 7,500 homes completely destroyed and approximately 20,000 that may need rebuilding, with safer sites being considered [6].

The scale and spread of these losses strengthen the Capital Sustainability finding. They do not prove that a specific L-ARC™ investment package would have avoided a known dollar amount. They do show that monitoring redundancy, protected communications, safe siting, drills, evacuation infrastructure, resilient energy systems and maintenance are recurring capital commitments whose value lies in reducing future exposure and preserving the ability to act before impact.

CSI dimension	Score	Evidence interpretation
C1 Resilience investment	30	Severe exposure of settlements, roads, hydropower and other corridor infrastructure indicates limited demonstrated physical resilience.
C2 Preparedness/capability	50	Substantial national search, rescue and coordination capacity was mobilised; two workers were rescued alive from the Trishuli 3A site on 4 September [1].
C3 Monitoring sustainability	30	Critical sensing/communications lacked sufficient fail-safe resilience to the event [5].
C4 Financial continuity/mobilisation	50	Government and international resources mobilised rapidly after impact; pre-impact resilience finance was less visible.
C5 Resilience value/durability	50	Reconstruction is shifting toward safer siting and warning-system resilience, but durability remains to be demonstrated [5,6].

Nepal CSI = 42/100 - Partially established

The central capital question is whether reconstruction merely restores assets or deliberately reduces future exposure while hardening sensing, communications, evacuation and critical infrastructure.

3.6 Nepal interface profile and maturity

Interface	Formal ratio status	Diagnosis	Evidence
IGC	National denominator not estimable	Weak-Moderate systemically; strong local example	Late/failed system-level warning continuity; school IGC 100% at n=1 [4,5]

Interface	Formal ratio status	Diagnosis	Evidence
GCR	Not estimable	Moderate for emergency mobilisation; weaker for pre-impact structural prevention	13,093 people rescued by 4 Sep; multinational tunnel/search support mobilised; major pre-impact infrastructure exposure remained [1-3].
CIR	Not yet estimable	Longitudinal test pending	Plans for seismic sensors, cameras, VSAT, real-time data and safer siting [5,6]

Maturity: Level 3 - Aware

Risk is monitored and institutionally understood, but intelligence, governance and resilient capital are not yet integrated sufficiently to manage the fastest cascading mountain hazards. Specific local systems may operate at higher maturity.

3.7 What Nepal validates

- Mechanism trace: warning -> recognised authority -> rapid decision -> reduced exposure before structural destruction.
- Diagnostic contrast: the same event contains a successful local hand-off and broader monitoring/warning discontinuity.
- Intelligence continuity: sensor presence is insufficient if the system does not detect sensor failure or retain redundant communication.
- Governance Legitimacy: authority and judgement matter most when certainty is incomplete and time is short.
- Capital Sustainability: the roughly US\$2.6 billion initial multi-sector loss estimate and post-event reconstruction choices show why resilient capital must reduce future exposure while hardening sensing, communications and critical infrastructure [1].
- Independent convergence: post-event plans for sensors, cameras, VSAT redundancy, real-time information and safer siting converge with the architecture L-ARC™ specifies, without implying adoption of L-ARC™.

COMPARATOR 1: WILDFIRE

4. Maui 2023 - Lahaina wildfire

The Maui case tests L-ARC™ against a fast-moving wildfire in a high-income jurisdiction. Maui County's Emergency Management Agency After-Action Report records 102 fatalities. Red Flag Warnings were issued on 6 August and the Emergency Operations Center was partially activated; the Lahaina fire reignited at 2:55 p.m. on 8 August and the EOC reached full activation at 4:30 p.m. [9]. FEMA/USFA reported more than 2,200 structures destroyed and approximately US\$5.5 billion in damage [11].

The Hawaii Attorney General/FSRI analysis found that only one siren from the All-Hazard Outdoor Warning Siren System was operable within the Lahaina burn perimeter on 8 August 2023 and that the siren system had not previously been used to warn of wildland-urban-interface fires [10]. These facts make the case particularly relevant to the Intelligence-to-Governance interface.

Observed fact	Value	L-ARC relevance
Fatalities	102 [9,12]	Outcome severity
Structures destroyed	>2,200 [11]	Exposure / capital resilience
Estimated damage	~US\$5.5 billion [11]	Resilience dividend / avoided-loss context
Red Flag warning	Issued 6 Aug; EOC partially activated [9]	Risk intelligence existed before impact
Operable siren in burn perimeter	1 [10]	Last-mile warning continuity
Index	Score	Retrospective interpretation
All	46	Fire-weather intelligence existed, but last-mile continuity and actionability were weak.
GLI	34	Decision, communication and evacuation conversion were fragmented during rapid escalation.
CSI	34	Communications, evacuation resilience and wildfire-preparedness capital were insufficient for event scale.
GPI	33	Formal authority existed, but effective permission/coordination under rapidly changing conditions was weak.

Provisional LCCS = 37.6/100. Formal system-wide IGC/GCR/CIR percentages are not estimable from the available denominator data.

Interface diagnosis

IGC: Weak. GCR: Weak-Moderate. CIR: Emerging post-event. The case supports the proposition that advance hazard information is not equivalent to protective intelligence unless it survives communications failure and reaches an actionable governance pathway.

Maturity: Level 3 - Aware

The wildfire threat was known and monitored. The end-to-end conversion from hazard intelligence to warning, evacuation and resilient infrastructure was not yet consistently anticipatory.

COMPARATOR 2: EXTREME HEAT

5. Europe 2022-2023 - extreme heat

Europe is an important counter-example to any theory that equates better forecasting with resilience. ECMWF reported that the June 2022 heatwaves were well predicted, and its July analysis noted that the event had been forecast at least three weeks in advance, albeit with uncertainty about severity [13]. Copernicus found summer 2022 to be Europe's hottest on record, 0.4°C warmer than the previous record summer [14].

WHO's July 2026 heat-and-health synthesis cites an estimated 61,672 heat-related excess deaths in Europe during summer 2022 [15]. WHO Europe reports more than 61,000 heat deaths in 2022 and approximately 47,500 in 2023 across 35 countries [16]. The EEA found that 21 of 38 EEA countries had heat-health action plans and 20 of 38 had heat-health surveillance systems - approximately 55.3% and 52.6%, respectively [17]. The coexistence of strong meteorological foresight with large mortality and uneven health-governance coverage is the key L-ARC™ signal: anticipation created time, but protection still depended on governance and durable adaptive capacity.

Observed fact	Value	L-ARC relevance
Forecast horizon	Event signal visible weeks ahead [13]	Strong anticipatory opportunity
Summer temperature record	0.4°C above previous record summer [14]	Hazard severity
Estimated heat deaths 2022	61,672 [15]	Outcome severity
Estimated heat deaths 2023	~47,000 [16]	Persistence across years
Countries with heat-health action plans	21/38 = 55.3% [17]	Governance/preparedness coverage
Countries with surveillance systems	20/38 = 52.6% [17]	Health intelligence coverage
Index	Score	Retrospective interpretation
All	82	Forecasting and monitoring were strong; meteorological anticipation was not the primary bottleneck.
GLI	58	Heat-health governance and surveillance coverage were substantial but uneven across jurisdictions.
CSI	54	Adaptation, health-system preparedness and urban resilience remained heterogeneous.
GPI	60	Many jurisdictions had formal authority and plans, but regional coverage was incomplete.

Provisional LCCS = 63.6/100.

Validation signal

Europe has the strongest All profile in the portfolio, yet severe health outcomes persisted. This is consistent with L-ARC™: intelligence creates opportunity, not guaranteed protection. Conversion through health governance, social protection, urban adaptation and sustained capital remains decisive.

Maturity: Level 3 - Aware at regional aggregate

Some national systems would qualify closer to Level 4 - Anticipatory. The regional aggregate deliberately preserves heterogeneity rather than awarding Europe a single uniform maturity level.

COMPARATOR 3: PANDEMIC

6. India 2021 - COVID-19 second wave

India tests L-ARC™ against a slow-onset biological hazard with long lead time relative to flash flood or wildfire. In November 2020, the Parliamentary Standing Committee on Health and Family Welfare noted that the threat of a second wave was high and that government hospital-bed capacity was inadequate for a significant rise in cases [18].

On 15 April 2021, with demand rising, the Government mapped oxygen sources for 12 high-burden states, identified another 100 hospitals for pressure-swing adsorption plants and decided to tender for import of 50,000 metric tonnes of medical oxygen [19]. A subsequent official update stated that 1,594 PSA plants were being established and described a geographic mismatch: roughly one-third of oxygen production was concentrated in eastern India while about 60% of demand was in the north and south [20]. The same update recorded liquid medical oxygen sales rising more than five-fold, from 1,559 MT/day on 31 March to more than 8,000 MT/day by 3 May, illustrating how rapidly logistics capacity had to scale [20].

WHO's South-East Asia situation report for 24-30 May recorded 1,364,668 new Indian cases in the week, a 26% fall from the previous week [21]. An official Japanese government reference citing Indian data recorded 27,894,800 cumulative cases and 325,972 cumulative deaths as of 30 May 2021 [22].

Observed fact	Value	L-ARC relevance
Second-wave warning	Threat assessed as high in Nov 2020 [18]	Advance risk knowledge
Government hospital beds	Committee judged inadequate [18]	Capital/preparedness constraint
Medical oxygen import decision	50,000 MT [19]	Emergency resource mobilisation
PSA plants being established	1,594 [20]	Capital-to-intelligence/capability reinforcement
New cases, week 21	1,364,668 [21]	Event scale
Cumulative cases / deaths	27,894,800 / 325,972 [22]	Outcome scale at cited snapshot
LMO demand acceleration	>5x: 1,559 MT/day (31 Mar) to >8,000 MT/day (3 May) [20]	Surge pace / logistics stress
Index	Score	Retrospective interpretation
All	50	Surveillance and expert warning existed; anticipation of surge scale and infrastructure needs was uneven.
GLI	54	Formal governance existed; conversion across national/state systems was uneven under extreme demand.
CSI	38	Distributed bed, oxygen-generation, storage and logistics resilience were insufficient at surge onset.
GPI	55	Authorities could act, but coordination and implementation lagged behind rapidly expanding demand.

Provisional LCCS = 46.8/100.

Interface diagnosis

IGC: Moderate. GCR: Weak-Moderate before and during surge onset. CIR: Stronger after the wave as oxygen-generation and distribution capacity expanded. The case illustrates that national aggregate supply can coexist with severe local access constraints when conversion and logistics are weak.

Maturity: Level 3 - Aware

The strategic need was recognised, but distributed operational capacity was not fully in place before peak demand. This is an institutional finding, not a partisan attribution of responsibility.

COMPARATOR 4: TROPICAL CYCLONE

7. Philippines 2021 - Super Typhoon Rai/Odette

The Philippines is deliberately included as a positive contrast. On 16 December 2021, the government reported Tropical Cyclone Warning Signal #4 in multiple areas, prepositioned standby funds and family food packs, and recorded 98,091 people already pre-emptively evacuated by 11:00 a.m. [23]. UNDRR later reported that more than 400,000 people were evacuated before impact and linked this capacity to strengthened early-warning and disaster-management systems developed after earlier typhoons [24].

The storm still caused catastrophic damage. NDRRMC final figures reproduced by the Asian Disaster Reduction Center record 405 deaths, 52 missing, 1,371 injured, 7,846,542 affected people and 2,028,894 houses totally or partially destroyed [25].

Observed fact	Value	L-ARC relevance
Pre-emptive evacuees by 11:00, 16 Dec	98,091 [23]	Documented early conversion before landfall
Pre-impact evacuation, later UNDRR account	>400,000 [24]	Strong IGC / governance conversion
Deaths	405 [25]	Outcome severity despite anticipatory action
Affected population	7,846,542 [25]	Scale and exposure
Houses damaged/destroyed	2,028,894 [25]	Capital and structural vulnerability
Standby resources before landfall	Funds and family food packs prepositioned [23]	GCR / resource readiness
Index	Score	Retrospective interpretation
All	74	Forecasting, warnings and risk communication were strong enough to support pre-emptive action.
GLI	70	Pre-agreed authority and local disaster governance converted warning into evacuation.
CSI	54	Emergency preparedness was strong, but structural exposure remained very high.
GPI	70	Decision rights and willingness to act before landfall were clearly demonstrated.

Provisional LCCS = 65.4/100.

Interface diagnosis

IGC: Strong. GCR: Strong for emergency mobilisation. CIR: Moderate-Strong, reflecting institutional learning after earlier typhoons. The case demonstrates the positive direction of the L-ARC™ mechanism: strong warning plus governance permission creates preventive opportunity, even when structural vulnerability remains high.

Maturity: Level 4 - Anticipatory

Early intelligence systematically informed pre-emptive decisions. The remaining gap was less about warning and more about the durability of buildings, infrastructure and exposure reduction.

COMPARATOR 5: CRITICAL INFRASTRUCTURE

8. Texas 2021 - electricity-grid failure

Texas provides the portfolio's clearest test of learning and capital reinforcement. FERC/NERC reported that ERCOT ordered 20,000 MW of rolling blackouts to prevent grid collapse and that more than 4.5 million people in Texas lost power [26]. The preliminary inquiry recorded 61,800 MW of generation lost during the freeze [27].

The final report found that freezing issues (44.2%) and fuel issues (31.4%) together accounted for 75.6% of unplanned generating-unit outages, derates and failures to start. It estimated that protecting just four categories of power-plant components from icing and freezing could have reduced ERCOT outages by 67% [26].

The event was not the first warning. FERC Chairman Rich Glick stated that a similar inquiry followed Texas extreme cold in 2011 but that its recommendations were not acted upon [27]. Texas DSHS subsequently confirmed 246 winter-storm-related deaths across 77 counties; 161 deaths (65.4%) were associated with extreme cold exposure [28].

Observed fact	Value	L-ARC relevance
Rolling blackouts ordered	20,000 MW [26]	System preservation under crisis
People losing power	>4.5 million [26]	Population exposure
Generation loss in inquiry	61,800 MW [27]	Infrastructure failure scale
Freezing + fuel contribution	75.6% [26]	Known technical vulnerability
Potential outage reduction from four component protections	67% in ERCOT [26]	Preventive capital opportunity
Confirmed storm-related deaths	246 across 77 counties [28]	Human impact
2011 inquiry	Recommendations not acted on [27]	CIR / institutional learning failure
Index	Score	Retrospective interpretation
All	50	Cold-weather vulnerability was known; risk information did not become sufficient preventive readiness.
GLI	38	Knowledge did not convert into strong enough binding governance and preparedness requirements.
CSI	38	Winterisation and gas-electric resilience investment were insufficient before recurrence.
GPI	35	Governance authority and incentives did not produce sufficiently decisive pre-event action.

Provisional LCCS = 41.6/100.

Interface diagnosis

IGC: Weak. GCR: Weak. CIR: Weak before 2021. Texas is a strong example of why recognising a vulnerability and repeatedly studying it are not equivalent to financing and enforcing prevention.

Maturity: Level 3 - Aware

The system was not ignorant of the hazard. It was aware without becoming adequately anticipatory.

CROSS-CASE RESULTS

9. Cross-case results

The six cases show that L-ARC™ does not merely explain low-capability systems. It differentiates high intelligence with incomplete conversion (Europe), relatively successful warning-to-action conversion (Philippines), local success within wider system discontinuity (Nepal), and recurring knowledge without sufficient capital reinforcement (Texas).




Figure 1. Retrospective All, GLI and CSI profiles. Scores are illustrative author-coded research outputs intended to test framework coherence, not official institutional ratings or definitive rankings.

Cross-case pattern

Across the six-case portfolio, the mean evidence-coded scores are All 57.3, GLI 52.0 and CSI 43.3. These are not population estimates or league-table ratings. They show a consistent portfolio pattern: institutions in these cases were generally better at recognising risk than at converting risk knowledge into durable governance and capital capability.

9.1 Interface profile across hazards

Case	IGC	GCR	CIR
Nepal	Weak-Moderate systemically; 100% school micro-case, n=1	Moderate response; weaker pre-event structural conversion	Not yet assessable; post-event plans visible
Maui	Weak	Weak-Moderate	Emerging post-event
Europe	Moderate	Moderate	Moderate
India	Moderate	Weak-Moderate pre-wave	Stronger post-wave
Philippines	Strong	Strong emergency mobilisation	Moderate-Strong
Texas	Weak	Weak	Weak pre-event

Formal system-wide percentages are intentionally withheld where denominator definitions are unavailable. This makes the analysis less superficially quantitative but more defensible.

9.2 Maturity profile

Level	Name	Definition	Cases in this paper
1	Fragile	Core capabilities largely absent or fragmented; no dependable preventive pathway.	None
2	Reactive	Some capability exists, but action occurs mainly after disruption or after preventive lead time is lost.	None as primary classification
3	Aware	Risk monitored and recognised; one or more critical conversion interfaces remain incomplete or inconsistent.	Nepal, Maui, Europe regional aggregate, India, Texas
4	Anticipatory	Early intelligence is systemically converted into timely, legitimate preventive action before impact.	Philippines
5	Adaptive	Anticipatory operation plus implemented learning that changes future protocols, thresholds or investment.	Not established
6	Resilient	Adaptive integration, redundancy and continuous improvement sustained across multiple cycles.	Not established

PROPOSITION TESTS

10. Test of the six L-ARC™ propositions

Proposition	Statement	Finding	Cross-case evidence
P1 Anticipatory Intelligence	Higher-quality anticipatory intelligence increases opportunity for timely prevention.	Supported	Philippines and Nepal school show protective opportunity; Europe shows opportunity is not sufficient by itself.
P2 Governance Conversion	Strong intelligence does not reliably produce prevention where Governance Legitimacy is weak.	Strongly supported	Maui, Texas and uneven European heat-health conversion.
P3 Capital Durability	Authorised intent does not produce durable resilience where Capital Sustainability is weak.	Strongly supported	Texas winterisation; India oxygen/logistics; Nepal infrastructure exposure.
P4 Interface Effect	Similar capability can yield different performance when interfaces differ.	Supported	Nepal within-event contrast; Philippines vs Maui warning conversion.
P5 Reinforcement	Investment that strengthens monitoring/learning improves subsequent preparedness.	Partially supported	Philippines post-Haiyan learning; India and Maui post-event reinforcement. Longitudinal causal test still required.
P6 Bottleneck	A severely weak capability or interface constrains the continuum.	Strongly supported diagnostically	High AI in Europe does not eliminate governance/capital gaps; Texas knowledge did not offset weak capital conversion.

Strongest result

The most consistent evidence is for P2, P3 and P6: resilience failure is repeatedly associated with breaks between intelligence, legitimate action and durable capability. P4 - interface effectiveness - is the most distinctive and now most testable L-ARC™ proposition.

UN AND INTERNATIONAL FRAMEWORK ALIGNMENT

11. Alignment with Sendai, Early Warnings for All and UNDRR 2026-2030

L-ARC™ should be positioned as complementary to established international frameworks, not as their replacement. The Sendai Framework's four priorities are understanding disaster risk; strengthening disaster-risk governance; investing in disaster-risk reduction for resilience; and enhancing preparedness and Build Back Better [29]. Early Warnings for All is organised around disaster-risk knowledge; detection/observation/monitoring/forecasting; warning dissemination and communication; and preparedness/response capabilities [30]. UNDRR's Strategic Framework 2026-2030 centres four catalytic areas: risk knowledge, locally led DRR, financing for DRR and recovery readiness, and explicitly seeks an open, interoperable and AI-enabled risk information ecosystem [31].

International architecture	Closest L-ARC™ correspondence	Potential incremental L-ARC™ contribution
Sendai Priority 1 - Understanding risk	All	Operationalises continuity, context, explainability and actionability of intelligence.
Sendai Priority 2 - Risk governance	GLI + GPI + IGC	Makes legitimacy of early action under uncertainty and conversion latency explicit.
Sendai Priority 3 - Investing in resilience	CSI + GCR + CIR	Separates post-event finance from persistent preventive capital and reinforcement of future intelligence.
Sendai Priority 4 - Preparedness / Build Back Better	Operating model + Adaptive Learning	Links outcome learning directly back into thresholds, monitoring and capital.
EW4All pillars	All + operating stages + GLI	Adds an explicit institutional permission and accountability lens at the warning-to-action hand-off.
UNDRR 2026-2030 catalytic areas	Risk knowledge + governance + financing + recovery readiness	Provides a recursive diagnostic across capability interfaces rather than a new global policy framework.

11.1 Relationship to AWARE

AWARE is the nearest peer-reviewed comparator because it explicitly links early warning, early action and early finance, with collaborative governance and feedback [32]. L-ARC™ does not claim first use of those ingredients. Its narrower architectural claim is that explainable anticipatory intelligence, legitimacy of early institutional action under uncertainty, and sustained preventive capital are named as co-dependent institutional capabilities, while the interfaces between them are treated as a primary failure surface.

Architectural differentiation

Among the major international and peer-reviewed frameworks examined, L-ARC™ is distinctive in making explainable AI-enabled Anticipatory Intelligence, Governance Legitimacy for early action under uncertainty, and Capital Sustainability the named co-dependent constructs of a recursive resilience continuum, with interface performance explicitly measured.

ROBUSTNESS, LIMITATIONS AND CLAIM BOUNDARIES

11.2 Institutional implications

The cross-hazard result is most useful when translated into institutional questions rather than treated as a ranking exercise. The practical value of L-ARC™ lies in identifying where a preventive-resilience chain is strong, where it is brittle, and which hand-off is most likely to consume the remaining lead time or prevent authorised action from becoming durable capability.

Institutional audience	L-ARC™ diagnostic question	Potential use
Government / DRR agencies	Where do warning, authority and escalation lose critical time?	Clarify triggers, escalation pathways, last-mile communication, redundancy and drills.
Critical infrastructure operators	Which sensing, communication or operating dependencies fail under stress?	Prioritise redundancy, continuity investment and pre-impact resource mobilisation.
Resilience finance / investment	Does capital merely restore assets, or does it reduce exposure and strengthen future intelligence?	Distinguish recovery spending from persistent preventive capital and reinforcement.
Researchers / assurance bodies	Can independent coders reproduce the capability and interface diagnosis?	Test reliability, construct validity, comparator performance and prospective utility.

Positioning guardrail: this white paper does not present L-ARC™ as a certification system or as a replacement for Sendai, Early Warnings for All or AWARE. Its proposed incremental value is diagnostic specificity at the interfaces where knowledge, permission, resources and long-term reinforcement must connect.

11.3 A practical L-ARC™ diagnostic cycle

1. Define the unit	Specify the hazard, geography, institution and decision boundary being assessed.
2. Build the evidence pack	Use a fixed chronology and evidence hierarchy; separate reported facts, derived arithmetic and interpretation.
3. Assess capabilities	Apply the pre-defined AI, GLI, CSI and GPI rules without changing constructs to fit the case.
4. Trace the interfaces	Examine Intelligence-to-Governance, Governance-to-Capital and Capital-to-Intelligence hand-offs; calculate formal ratios only where denominators are defensible.
5. Convert diagnosis into learning	Identify preventive changes, assign ownership, and test whether implemented investment improves future sensing, permission, mobilisation and continuity.

12. Robustness review: questions a serious institution should test

Challenge	Response
“This is just early warning.”	Partly. Nepal and the Philippines validate principles shared with EW4All and anticipatory action. L-ARC™ adds a specific institutional conversion lens: whether intelligence becomes legitimate action, whether action mobilises capacity, and whether capital preserves that ability over time.
“Sendai already has risk, governance and investment.”	Correct. Sendai is a major conceptual ancestor. L-ARC™ should be tested for incremental diagnostic specificity, not marketed as replacing Sendai.
“AWARE already links warning, action and finance.”	Correct. AWARE is the closest peer-reviewed comparator. The appropriate next step is direct coding of the same cases using both frameworks.
“The scores are subjective.”	The present scores are author-generated. They are illustrative research outputs intended to test framework coherence rather than establish definitive institutional rankings. The disclosed evidence bands and full score rationale make the judgement traceable; Appendix C specifies the independent reliability design needed to test reproducibility.
“The Nepal school result is heroic leadership, not governance.”	Leadership mattered. L-ARC™ asks whether an actor had recognised authority, trusted the signal and could mobilise resources. Replication should test the same mechanism in more formal decision protocols.

Challenge	Response
“The author is validating his own framework.”	Yes. The conflict is explicit. Safeguards are source hierarchy, fixed cut-offs, transparent arithmetic, competing frameworks, negative findings and invitation to independent replication.
“One portfolio cannot establish universal efficacy.”	Agreed. The paper claims cross-hazard retrospective mechanism validation, not a universal causal treatment effect.

12.1 What the evidence establishes - and what remains open

Claim	Status	Evidence boundary
Core Intelligence -> Governance mechanism	SUPPORTED AT MECHANISM LEVEL	Observed Nepal school sequence; positive Philippines pre-emptive evacuation; contrasting Maui/Texas conversion failures.
Decision latency as a resilience variable	STRONGLY SUPPORTED	Nepal 14-minute protective lead time and 41-minute public-alert interval; not causally equivalent but operationally informative.
All/GLI/CSI applicability across hazard classes	SUPPORTED ACROSS SIX CASES	Same constructs remain interpretable across all six cases.
Interface diagnostic value	STRONGLY SUPPORTED	Repeated failure or success at IGC/GCR/CIR hand-offs.
Exact lives or losses avoided under L-ARC™	NOT ESTABLISHED	Requires counterfactual exposure and intervention-effect data.
Universal index calibration	NOT ESTABLISHED	Requires construct validation, empirical distributions and independent raters.
Superiority over Sendai/EW4All/AWARE	NOT ESTABLISHED	Requires direct comparative studies on common cases.
Prospective causal efficacy	NOT ESTABLISHED	Requires pre-registered field deployment or quasi-experimental design.

12.2 Potential falsifiers and disconfirming evidence

A defensible framework must specify what evidence would count against it. The following patterns would weaken, narrow or falsify specific L-ARC™ propositions rather than being reinterpreted after the fact.

Potential disconfirming pattern	Implication for L-ARC™
Repeated high resilience under comparable hazard exposure despite weak AI, GLI and CSI and weak interfaces.	Would challenge the Bottleneck Proposition and the claim that continuity across the three capabilities is materially important.
Protective outcomes repeatedly occur where credible intelligence never reaches a governance decision pathway before action.	Would challenge the Governance Conversion proposition or require a narrower account of what constitutes anticipatory intelligence.
Matched cases show no meaningful relationship between interface quality and differences in preventive performance.	Would weaken the Interface Effect and reduce the case for treating IGC/GCR/CIR as distinct diagnostic surfaces.
Longitudinal cases show durable resilience without persistent capital, monitoring redundancy or reinforcement after prior events.	Would weaken the Capital Durability and Reinforcement propositions.
Independent matched-case comparison finds AWARE, EW4AI or Sendai-derived coding explains the same bottlenecks equally well, with no reproducible incremental specificity from L-ARC™.	Would require the incremental-value claim to be narrowed; L-ARC™ could remain a synthesis but not claim added diagnostic specificity.
Prospective pilots repeatedly fail to identify actual bottlenecks before events, or targeted changes to predicted bottlenecks do not improve pre-specified process outcomes.	Would challenge prospective operational utility and any future causal-efficacy claim.

12.3 Alternative explanations

The patterns identified through L-ARC™ may also be partly explained by broader institutional variables not directly measured in this study. These include overall state capacity, economic development, infrastructure quality, disaster-management maturity, governance culture, and hazard severity. The present analysis does not claim that L-ARC™ is the sole explanation for observed outcomes. Rather, it tests whether Intelligence-to-Governance-to-Capital conversion provides a useful and repeatable diagnostic lens across heterogeneous hazards. Future matched-case and comparator-framework studies should assess the relative contribution of

INSTITUTIONAL ADOPTION AND INDEPENDENT REPLICATION

these alternative explanatory factors.

13. Pathway to institutional adoption and independent replication

The present six-case study provides an empirical base for institutional use and independent replication. A global institution can apply the same constructs and evidence rules to additional cases, compare results with established approaches, and use prospective pilots to test whether identified bottlenecks are visible before an event. Replication should strengthen or challenge the evidence without changing the core constructs to fit preferred outcomes.

Step	Institutional validation requirement
1. Maintain a stable Master Specification	Do not add foundational constructs solely to fit new cases.
2. Publish a Measurement and Indicator Protocol	Define every indicator, denominator, evidence requirement, normalisation rule and missing-data treatment.
3. Pre-specify the cross-hazard coding protocol	Pre-specify the scoring rule before adding further cases.

Step	Institutional validation requirement
4. Expand the case corpus	Target at least 20-30 heterogeneous successes, near misses and failures across hazards, regions and income groups.
5. Use independent blinded raters	At least two or three researchers not involved in framework development should code the same events.
6. Test inter-rater reliability	Use weighted kappa or intraclass correlation where appropriate.
7. Test construct validity	Assess whether All, GLI and CSI are empirically distinct from generic preparedness.
8. Conduct comparator studies	Code common cases using L-ARC™, AWARE, EW4All and relevant sector instruments.
9. Run prospective pilots	Specify anticipated bottlenecks before a hazard season or operational deployment.
10. Enable independent replication	Release a documented protocol and evidence register so universities and institutions can reproduce or challenge results.

Next credibility milestone

Independent coding with inter-rater reliability and matched-case comparison against AWARE, EW4All and Sendai-derived assessment should precede claims of comparative advantage. Additional case expansion is useful, but it is no longer the highest-value next step.

13.1 Priority publication design: matched comparator coding

The strongest next publication should shift the research question from whether L-ARC™ can explain disasters to whether it identifies reproducible institutional information that established frameworks do not capture as clearly. The same events, evidence windows and source packs should therefore be coded using L-ARC™, AWARE, Early Warnings for All and a Sendai-derived assessment rubric [29-32].

Design control	Pre-specified rule
Evidence parity	All frameworks receive the same source pack, chronology and unit of assessment.
Case parity	The same cases are used unless a framework is demonstrably outside scope; exclusions are documented before coding.
Framework fidelity	Each comparator is represented according to its authoritative published architecture; no strawman simplification.
Coder independence	Where feasible, independent researchers familiar with each framework code cases without being told an expected L-ARC™ result.
Reliability	Agreement is reported for L-ARC™ and, where methodologically possible, for comparator-derived coding rather than only comparing mean scores.
Comparison criteria	Bottleneck specificity, interface visibility, evidence traceability, actionability, reproducibility and conceptual redundancy are assessed using pre-specified criteria.
No automatic superiority	If comparators explain the same cases equally well and L-ARC™ adds no reliable diagnostic specificity, the incremental-value claim is narrowed.

This comparative design directly tests the strongest publication question: not merely whether L-ARC™ is coherent, but whether its explicit Intelligence -> Governance -> Capital interfaces add measurable diagnostic value beyond existing frameworks.

Statistical construct validation after reliability

After inter-rater reliability is established and the case corpus is sufficiently large, index development should proceed to formal construct validation. Exploratory and, where sample size permits, confirmatory factor analysis should test whether indicators cluster as hypothesised into Anticipatory Intelligence, Governance Legitimacy and Capital Sustainability. Convergent validity should test whether related L-ARC™ dimensions associate with established preparedness, early-warning, governance or resilience measures in the expected direction. Discriminant validity should test whether All, GLI and CSI remain empirically distinguishable rather than collapsing into a single generic preparedness construct. Criterion validity should test association with pre-

specified operational criteria such as warning-to-governance latency, evacuation or protective-action coverage, resource mobilisation, continuity failures and implemented resilience investment. These analyses require a substantially larger sample than the present six cases and should follow, not substitute for, reliability testing and matched-framework comparison.

13.2 Subsequent case-expansion portfolio

- Flood: additional South Asia and Latin America cases with both successful and failed warning conversion.
- Wildfire: Australian and Mediterranean cases, including prevention and near misses, not only catastrophic fires.
- Heat: city-level heat-health action plans with mortality surveillance.
- Cyclone: comparative cases with similar forecast quality but different evacuation and loss outcomes.
- Infrastructure: transport, water, digital and energy failures where prior warnings can be traced.
- Pandemic/health emergency: hospital surge, oxygen, supply chain and public-health decision pathways.
- Ecological crises: biodiversity and ecosystem thresholds where intervention legitimacy and capital durability are central.

CONCLUSION

14. Conclusion - evidence consistent with a cross-hazard resilience mechanism

The six cases are physically incomparable but institutionally comparable. Nepal's flood, Maui's wildfire, European heat, India's pandemic surge, Typhoon Rai and the Texas grid failure involve different hazards, lead times and governance systems. Yet each can be analysed through the same L-ARC™ question: did institutions convert anticipatory intelligence into legitimate action, action into resourced capability, and resilience capital into stronger future intelligence? The recurrence of this pattern across all six cases is the core empirical result.

Nepal is the most compelling anchor because the same event contains both sides of the continuum. At one school, approximately 14 minutes mattered: information reached a legitimate decision-maker, normal activity stopped, buses were redirected and more than 900 students plus 16 staff moved before the school was destroyed [4]. In the broader corridor, at least four gauges were lost and a public alert followed the border impact by about 41 minutes [5]. The comparison is not a national counterfactual; it is direct evidence that warning continuity, decision rights and residual lead time were operational variables in the same disaster.

The Philippines shows the mechanism in the positive direction: strong warning, pre-positioned resources and institutional permission supported mass pre-emptive evacuation [23,24]. Texas shows the reverse: prior knowledge and inquiry did not become sufficient winterisation and cross-system resilience before a recurrent event [26,27]. Europe, Maui and India provide intermediate patterns in which strong or partial intelligence was constrained by governance, last-mile conversion, logistics or capital durability. Taken together, these cases provide evidence consistent with the diagnostic relevance of the L-ARC™ continuum across hazard classes.

Principal study finding

On the evidence available at the stated cut-offs, the six cases provide evidence consistent with cross-hazard retrospective mechanism coherence and cross-geographic diagnostic utility for L-ARC™. The strongest support concerns Governance Conversion, Capital Durability, Interface Effectiveness and the Bottleneck Proposition. Universal calibration, comparative superiority and prospective causal effect sizes remain open empirical questions rather than assumptions.

14.1 Institutional positioning statement

Institutional positioning

L-ARC™ is a theoretically specified preventive-resilience framework with an empirically grounded retrospective evidence base across six materially different hazard classes. It operationalises Anticipatory Intelligence, Governance Legitimacy and Capital Sustainability and makes the interfaces between them directly testable. Nepal supplies a documented within-event warning-decision-action trace, while the comparator cases provide positive and negative contrasts. Numerical scores remain illustrative author-coded research outputs pending independent reliability testing and broader calibration.

14.2 Final research statement

L-ARC™ does not ask only: Did we see the risk? It asks: Did we understand it early enough? Could somebody legitimately act? Could the decision mobilise resources? Did action occur while prevention was still possible? Did investment make the next warning system stronger?

Across the six cases examined here, those questions consistently distinguish protective conversion, partial conversion and institutional failure. That is the central empirical result of this cross-hazard validation and Nepal case study.

White Paper Takeaways

What the six-case evidence allows an institutional reader to take forward

1	L-ARC™ survives a cross-hazard retrospective stress test at the mechanism and diagnostic level: the same Intelligence -> Governance -> Capital logic remains interpretable across all six cases.
2	The most distinctive empirical signal is conversion at the interfaces. Risk knowledge is necessary but not sufficient; protective value depends on timely governance disposition, resource mobilisation and reinforcement.
3	Decision latency is a resilience variable. Nepal shows that minutes, warning continuity and recognised authority can materially shape the opportunity to act, without supporting a claim about the exact national losses that were preventable.
4	High-quality anticipation does not guarantee resilience. Europe demonstrates that strong forecasting can coexist with major consequences when governance coverage and adaptive capacity are uneven.
5	Capital Sustainability is preventive, not merely financial. Monitoring redundancy, communications, infrastructure hardening and learning must remain operational across cycles rather than appear only after disaster.
6	The next credibility milestone is independent replication: blinded coding, inter-rater reliability, matched comparison against AWARE/EW4All/Sendai-derived rubrics, and prospective pilots.

Suggested citation

Lahiri, S. (2026). L-ARC™ Framework: Empirical Cross-Hazard Validation - Nepal 2026 Case Study and Comparative Evidence Across Six Hazards. Research White Paper, Version 2.4.

Use and interpretation

This white paper is intended for research, policy, institutional learning and framework evaluation. Scores are transparent author-coded research outputs, not official ratings, certification outcomes or disaster probabilities. Readers are encouraged to reproduce, challenge and compare the analysis using the disclosed evidence rules and source base.

REFERENCES

References

- [1] Government of Nepal, Ministry of Foreign Affairs. (2026, 4 September). Daily Situation Update on the Bhote Koshi River Floods - 04 September 2026 (updated 14:00 NST). <https://www.mofa.gov.np/content/1876/daily-situation-update-on-the-bhote-koshi/>
- [2] United Nations Development Programme. (2026). Nepal Floods 2026: RAPIDA key findings. <https://www.undp.org/asia-pacific/nepal-floods-2026>
- [3] United Nations Development Programme. (2026, 1 September). Nepal floods hit vulnerable communities as UNDP estimates 2.2 million tonnes of debris. <https://www.undp.org/asia-pacific/press-releases/nepal-floods-hit-vulnerable-communities-undp-estimates-22-million-tonnes-debris>
- [4] Reuters. (2026, 31 August). Last-minute flood warning saved over 900 Nepal students, school head says. <https://www.reuters.com/business/environment/last-minute-flood-warning-saved-over-900-nepal-students-school-head-says-2026-08-31/>
- [5] Reuters. (2026, 2 September). Nepal to rebuild early-warning system on China border after deadly flood, sources say. <https://www.reuters.com/business/environment/nepal-rebuild-early-warning-system-china-border-after-deadly-flood-sources-say-2026-09-02/>
- [6] Reuters. (2026, 3 September). Nepal must rebuild thousands of homes, no letup in search for missing. <https://www.reuters.com/business/environment/nepal-must-rebuild-thousands-homes-no-let-up-search-missing-2026-09-03/>
- [7] European Space Agency. (2026, 31 August). Nepal flash flood imaged by satellites. https://www.esa.int/Applications/Observing_the_Earth/Copernicus/Sentinel-2/Nepal_flash_flood_imaged_by_satellites
- [8] United Nations Office for Disaster Risk Reduction. (2026, 28 August). Statement on the disaster in Nepal and China by SRSG Kamal Kishore. <https://www.undrr.org/news/statement-disaster-nepal-and-china-kamal-kishore-special-representative-un-secretary-general>
- [9] Maui Emergency Management Agency. (2024). Maui Wildfires 2023 After-Action Report. <https://www.mauicounty.gov/DocumentCenter/View/151355/MEMA-2023-Wildfire-After-Action-Report>
- [10] State of Hawaii Department of the Attorney General / Fire Safety Research Institute. (2024). Lahaina Fire Incident Analysis - Findings and Recommendations. https://ag.hawaii.gov/wp-content/uploads/2024/09/FSRI_Lahaina_Report_Findings_and_Recommendations_09102024.pdf
- [11] U.S. Fire Administration / FEMA. (2024, 8 February). Preliminary After-Action Report: 2023 Maui Wildfire. <https://www.usfa.fema.gov/blog/preliminary-after-action-report-2023-maui-wildfire/>
- [12] Office of the Governor, State of Hawaii. (2024, 9 August). Healing and Hope: One Year After the Maui Wildfires. <https://governor.hawaii.gov/newsroom/office-of-the-governor-news-release-healing-and-hope-one-year-after-the-maui-wildfires/>
- [13] European Centre for Medium-Range Weather Forecasts. (2022, 11 July; 18 July). European heatwaves in June 2022; Update on European heatwave of July 2022. <https://www.ecmwf.int/>
- [14] Copernicus Climate Change Service. (2022). Summer 2022 Europe's hottest on record. <https://climate.copernicus.eu/copernicus-summer-2022-europes-hottest-record>
- [15] World Health Organization. (2026, 31 July). Heat and health. <https://www.who.int/news-room/fact-sheets/detail/climate-change-heat-and-health>
- [16] World Health Organization Regional Office for Europe. (2025, 30 May). Climate change and health. <https://www.who.int/europe/news-room/fact-sheets/item/climate-change>
- [17] European Environment Agency. (2024). The impacts of heat on health: surveillance and preparedness in Europe. <https://www.eea.europa.eu/en/analysis/publications/the-impacts-of-heat-on-health>
- [18] PRS Legislative Research. (2020). Outbreak of pandemic COVID-19 and its management - Standing Committee report summary. <https://prsindia.org/policy/report-summaries/outbreak-of-pandemic-covid-19-and-its-management>
- [19] Government of India, Press Information Bureau. (2021, 15 April). Update on Medical Oxygen Availability. <https://www.pib.gov.in/Pressreleaseshare.aspx?PRID=1712123&lang=2®=48>
- [20] Government of India, Press Information Bureau. (2021). Centre undertakes multiple initiatives to enhance oxygen availability, distribution and storage infrastructure. <https://www.pib.gov.in/Pressreleaseshare.aspx?PRID=1717459&lang=2®=48>
- [21] World Health Organization Regional Office for South-East Asia. (2021, 4 June). COVID-19 Weekly Situation Report, Week 21 (24-30 May 2021). <https://cdn.who.int/media/docs/default-source/searo/whe/coronavirus19/sear-weekly-reports/searo-weekly-situation-report-21-2021-.pdf>
- [22] Ministry of Foreign Affairs of Japan. (2021). Emergency Grant Aid in Response to COVID-19 in India - reference data on infections. https://www.mofa.go.jp/press/release/press1e_000198.html
- [23] Presidential Communications Office, Philippines. (2021, 16 December). On Typhoon Odette. <https://pco.gov.ph/PCO-content/on-typhoon-odette/>
- [24] United Nations Office for Disaster Risk Reduction. (2023). The Philippines: Ten years after Haiyan - Building Back Better. <https://www.undrr.org/news/ten-years-after-haiyan-building-back-better-philippines>
- [25] Asian Disaster Reduction Center. (2022). Philippines: Typhoon Rai/Odette, NDRRMC final disaster information. https://www.adrc.asia/view_disaster_en.php?Key=2507&Lang=en
- [26] Federal Energy Regulatory Commission / North American Electric Reliability Corporation. (2021, 16 November). Final Report on February 2021 Freeze Underscores Winterization Recommendations. <https://www.ferc.gov/news-events/news/final-report-february-2021-freeze-underscores-winterization-recommendations>
- [27] Federal Energy Regulatory Commission / North American Electric Reliability Corporation. (2021, 23 September). FERC, NERC Staff Review 2021 Winter Freeze, Recommend Standards Improvements. <https://www.ferc.gov/news-events/news/ferc-nerc-staff-review-2021-winter-freeze-recommend-standards-improvements>

- [28] Texas Department of State Health Services. (2021, 31 December). February 2021 Winter Storm-Related Deaths - Texas. https://www.dshs.texas.gov/sites/default/files/news/updates/SMOC_FebWinterStorm_MortalitySurvReport_12-30-21.pdf
- [29] United Nations Office for Disaster Risk Reduction. (2015). Sendai Framework for Disaster Risk Reduction 2015-2030. <https://www.undrr.org/publication/sendai-framework-disaster-risk-reduction-2015-2030>
- [30] World Meteorological Organization. (2026). Early Warnings for All. <https://wmo.int/all-activities/build-resilience/early-warnings-all>
- [31] United Nations Office for Disaster Risk Reduction. (2026). UNDRR Strategic Framework 2026-2030. <https://www.undrr.org/strategic-framework-2026-2030>
- [32] Attah, E. M. N. A. N., & Amarnath, G. (2025). A framework for addressing the interconnectedness of early warning to action and finance to strengthen multiscale institutional responses to climate shocks and disasters. *Climate Risk Management*, 47, 100689. <https://doi.org/10.1016/j.crm.2025.100689>
- [33] Lahiri, S. (2026). **L-ARC™ Master Framework, Operating Model, Measurement and Validation Specification v2.0. Author framework specification.**
- [34] McHugh, M. L. (2012). Interrater reliability: the kappa statistic. *Biochemia Medica*, 22(3), 276-282. <https://doi.org/10.11613/BM.2012.031>
- [35] Koo, T. K., & Li, M. Y. (2016). A guideline of selecting and reporting intraclass correlation coefficients for reliability research. *Journal of Chiropractic Medicine*, 15(2), 155-163. <https://doi.org/10.1016/j.jcm.2016.02.012>

APPENDIX A - CROSS-HAZARD EVIDENCE-TO-CLAIM LEDGER

Appendix A. Evidence-to-claim ledger

Evidence	Type	L-ARC test	Permitted inference	Source
Nepal school: ~14-minute warning; >900 students +16 staff evacuated; school destroyed	Direct observed sequence	IGC / GLI	Mechanism validation; does not quantify national lives saved	[4]
Nepal border impact 08:32; public alert 09:13	Contemporaneous reporting	IGC / decision latency	Establishes ~41-minute interval; not proof that delay caused all losses	[5]
Nepal: at least four gauges lost; monitoring not extensive	Official-source interviews via Reuters	All continuity	Sensor presence did not equal resilient intelligence	[5]
Nepal: 84,270 affected; 76% in above-average vulnerability municipalities	UNDP/government data	Exposure / CSI	Derived ~64,045 is arithmetic; vulnerability classification is municipality-level	[2]
Nepal: initial total losses NPR 387.54bn (~US\$2.6bn); immediate early-recovery need NPR 7.95bn (~US\$53m)	Government / NDRRMA initial assessment	CSI / GCR	Demonstrates cross-sector capital exposure and near-term recovery requirement; figures are initial and not a final comprehensive loss assessment	[1]
Maui: Red Flag warnings; 1 operable siren in burn perimeter	Official after-action/investigation	All / IGC	Advance hazard knowledge existed; last-mile warning continuity was weak	[9,10]
Europe: heatwave well predicted; 61,672 estimated heat deaths in 2022	ECMWF / WHO	All vs GLI/CSI	High All does not guarantee protection	[13,15]
India: second-wave threat high; hospital beds inadequate; later 50,000 MT oxygen import decision	Parliamentary summary / PIB	All / GCR / CSI	Known risk did not equal distributed surge capacity	[18,19]
Philippines: >400,000 pre-emptively evacuated before Rai	UNDRR	IGC / GPI	Positive conversion evidence	[24]
Texas: 2011 recommendations not acted upon; 2021 freeze causes 20,000 MW load shed	FERC/NERC	CIR / GCR	Learning and capital reinforcement failure	[26,27]

APPENDIX B - FULL RETROSPECTIVE SCORE RATIONALE

Appendix B. Full score rationale

The following dimension values use only the five midpoint coding values 10, 30, 50, 70 and 90. Index scores are arithmetic means. These scores are illustrative author-coded research outputs intended to test framework coherence, not definitive institutional rankings. They are shown in full so independent coders can reproduce, challenge or revise the ratings using the same evidence rules.

Case	A1-A5	All	G1-G5	GLI	C1-C5	CSI
Nepal flood 2026	50/30/30/50/50	42	70/50/70/50/50	58	30/50/30/50/50	42
Maui wildfire 2023	70/30/50/50/30	46	50/30/30/30/30	34	30/30/30/50/30	34
Europe heat 2022-23	90/90/90/70/70	82	70/70/50/50/50	58	50/50/70/50/50	54
India COVID-19 2021	70/50/50/50/30	50	70/50/50/50/50	54	30/50/30/50/30	38
Philippines Rai 2021	90/70/70/70/70	74	70/70/70/70/70	70	50/70/50/50/50	54
Texas grid 2021	70/50/50/50/30	50	50/30/30/50/30	38	30/30/30/50/50	38

GPI values are separate governance-permission diagnostics derived from authority, escalation, threshold, confidence and defensibility evidence; they are not added to LCCS and therefore do not double-count Governance Legitimacy.

APPENDIX C - PLANNED RELIABILITY VALIDATION

Appendix C. Planned Reliability Validation

The present scores were generated by the framework creator and therefore do not yet establish inter-rater reproducibility. This appendix pre-specifies the reliability design that should be used before L-ARC™ scores are treated as externally calibrated institutional measurements. The objective is not to force agreement, but to determine whether independent researchers applying the same evidence rules reach sufficiently consistent judgements.

Reliability statistic	Planned use in L-ARC™	Reporting requirement
Cohen's kappa (κ)	Binary or nominal judgements, such as indicator eligibility, presence/absence of qualifying evidence, or categorical evidence decisions made by two independent coders.	Report κ , raw agreement, 95% confidence interval where estimable, category prevalence and disagreement counts [34].
Weighted kappa (κ_w)	Ordinal evidence-band decisions, especially the five ordered coding anchors 10/30/50/70/90. The weighting scheme must be selected before case results are inspected.	Report κ_w with its weighting rule, confidence interval, item-level disagreement and sensitivity to alternative defensible weighting where material [34].
Intraclass Correlation Coefficient (ICC)	Index-level outputs such as AI, GLI, CSI, GPI and LCCS when treated as continuous research measurements across independent raters.	Use a two-way random-effects, absolute-agreement model when raters are intended to generalise beyond the study; report single-measure ICC and, if an average of raters is used, average-measure ICC with 95% confidence intervals [35].

Planned coding procedure

- Use at least two, preferably three, independent coders who were not involved in creating L-ARC™.
- Provide each coder with the same fixed evidence pack, case chronology, unit of assessment and coder manual, without expected L-ARC™ scores.
- Calibrate on 4-6 heterogeneous cases before full-corpus coding; record ambiguous wording rather than resolving disagreements by instruction from the framework creator.
- After any protocol clarification, repeat calibration and then keep the revised manual stable for the study.
- Double-code at least 25-30% of a larger case corpus to detect coder drift and report both agreement and disagreement.
- Report reliability at indicator, dimension, index and maturity-classification levels; do not rely on a single summary coefficient.

Pre-specified interpretation

As practical study targets rather than universal scientific laws, weighted kappa of approximately 0.70 or higher for ordinal coding and ICC of approximately 0.75 or higher for index outputs can be used as initial quality gates, with confidence intervals and disagreement patterns reported alongside the point estimate [34,35]. Thresholds must be specified before results are known and should not replace substantive examination of where coders disagree.

Failure rule

If reliability falls below a pre-specified target, disagreement must not be averaged away. The ambiguous construct, evidence rule or coding instruction should be identified, revised transparently, re-tested on calibration cases and documented before further inference. A reliability failure is evidence about the measurement protocol and should be reported as such.

APPENDIX D - EVIDENCE CONFIDENCE SCALE

Appendix D. Evidence Confidence Scale

Evidence confidence is reported separately from All, GLI, CSI, GPI and LCCS. It describes the strength, completeness and traceability of the evidence available for a case; it is not a probability that the author-assigned score is correct and it is not mathematically folded into any L-ARC™ index. When evidence meets characteristics spanning two levels, the lower level is assigned and the limiting condition is stated.

Confidence level	Minimum evidence characteristics	Typical limiting condition
High	Multiple independent Tier 1 sources, including an official investigation, final official dataset or equivalent authoritative record where available; key chronology and denominators are substantially complete; no unresolved contradiction material to the primary inference.	Minor residual uncertainty does not materially change the case-level mechanism interpretation.
Moderate-High	At least one Tier 1 source plus independent Tier 1/Tier 2 convergence; primary chronology and mechanism claims are traceable; evidence is sufficiently complete for the stated inference.	Some data remain evolving, a denominator is incomplete, or a specific timeline relies on corroborated Tier 3 reporting.
Moderate	Credible Tier 1/Tier 2 evidence supports the main interpretation, but the unit of assessment is heterogeneous, denominator data are materially incomplete, evidence is aggregated across jurisdictions, or comparability is constrained.	The mechanism can be assessed, but confidence in a single case-wide quantitative profile is limited.
Low / insufficient	Evidence is sparse, substantially single-source, materially contradictory, or key claims depend on uncorroborated Tier 3/self-reported material.	Not sufficient for a strong cross-case inference; scoring should be withheld where the minimum evidence threshold is not met.

Application to the current portfolio

Nepal is classified Moderate-High because Government of Nepal, UNDP, UNDRR and ESA evidence converges with contemporaneous event-timeline reporting, while casualty/loss reporting was still evolving and national interface denominators remain incomplete. Europe is classified Moderate because forecast and health evidence are authoritative, but the unit of assessment is a heterogeneous regional aggregate with uneven national heat-health planning and surveillance coverage. The same rubric is applied to the remaining cases according to source authority, convergence, denominator completeness, temporal fit and traceability.