Rain to Phone

What an open, end-to-end flood monitoring system shows about Thailand's coordination gap

Poon Thiengburanathum[1,2] and Non Arkaraprasertkul[3,*]

[1]Faculty of Engineering, Chiang Mai University, Thailand
[2]Research and Innovation Acceleration Agency for Competitiveness and Area Development (RCAD), Thailand
[3]Digital Economy Promotion Agency (depa), Thailand

[*]Correspondence through https://flood.nonarkara.org

Abstract. On 19 August 2026 the river in Nan province reached its banks about four hours after the rain began, and the first reported disaster-agency cell broadcast came about four hours after that. Two people died and more than 16,000 were affected. Every number needed to see it coming was already public. We ask whether Thailand's long record of flood failure comes from three gaps: agencies do not share data in a form others can use, warnings do not respond to the data that exist, and no single chain runs from rain gauge to drain to phone. We ask whether an independent system can close the gaps without waiting for institutions to merge. We built one, FloodDash, from 44 open feeds, and report its first 124 days of operation (3 July to 4 November 2026) from its own logs. It raised 13,766 alerts. The median alert followed its gauge reading by 6.7 minutes. The system was down for 47 of 1,090 observed hours. Against the gauge network it reads, its 48-hour watchlist caught 84% of bank overflows with a median lead of 8.2 hours, at 12.5% precision. Rain gauges of two agencies within 2 km of each other correlated at only $r = 0.67$. Of 1,479 water-level gauges, 36 carry bank levels that cannot be right and only 52% reported in the previous day. The last mile failed: 29 people subscribed. Integration and speed can be built without permission. Reach cannot. We propose a minimum open interface for public hydrometeorological data and a verified relay between independent systems and the official warning channel.

Keywords: flood early warning; inter-agency coordination; open data; Thailand; last-mile warning; hydroinformatics

1 Introduction

At 01:00 on 19 August 2026, rain over Nan province passed 50 mm h^{-1} at the heaviest gauge. By the 04:00 hour three river gauges had reached their bank levels. A cell broadcast from the Department of Disaster Prevention and Mitigation (DDPM) to Pua district was reported at 08:30. Two people died and more than 16,000 were affected (The Nation, 2026; Thairath, 2026). Figure 1 shows who knew what, and when.

Nothing in that timeline was secret. The gauges were reporting. The numbers sat on public platforms. What did not exist was one place that read all of them at once and said something to a person.

This paper is about that place. It asks one question, and the question has three parts.

Thailand's record on floods is bad enough to need no adjectives. The 2011 floods covered 66 of 77 provinces, affected more than 13 million people and caused more than 680 deaths and THB 1.43 trillion (USD 46.5 billion) of damage and loss (World Bank, 2012). In November 2025 flooding in the south killed at least 145 people, 110 of them in Songkhla, where Hat Yai sits (Al Jazeera, 2025). Between those two floods, Thailand wrote a Water Resources Act and created a national office to coordinate water agencies (Royal Thai Government, 2018). The floods kept coming.

We take three explanations seriously, because they are the ones we hear most from people who work on the problem. *H1, coordination:* the agencies that measure, forecast and respond do not share data in a form the others can use. *H2, responsiveness:* the data that do exist do not drive action fast enough. *H3, the chain:* no system runs end to end, from the rain gauge through the river and the drain to a phone in someone's hand.

Our research question follows. *Can an independent system tackle all three, bypassing bureaucracy and giving people something they can use right away?*

We answer by building it and then measuring it. FloodDash is an open, continuously running system. It reads 44 public feeds, judges them, and sends alerts to phones through Telegram, Web Push, a live event stream and a text-only page. The second author built it, with AI coding assistants, between 3 July and 4 November 2026. It runs on one computer. We report what its logs show.

The answer is mixed, and we say so early because the mixed part is the finding. The system closes much of H2 and H3 on the data side. It cannot close the last metre of H3, the part where a warning must reach a stranger and be believed, and it can only describe H1, not repair it. Section 5 says what follows from that.

2 Background

2.1 What the two floods showed

The World Bank's assessment of 2011 is blunt about sequence. The floods ran from July to December. The government set up centralised flood monitoring and relief operations in mid-August, after they had begun (World Bank, 2012). It also names the cause as continuous monsoon rain and numerous dam breaches, and finds that

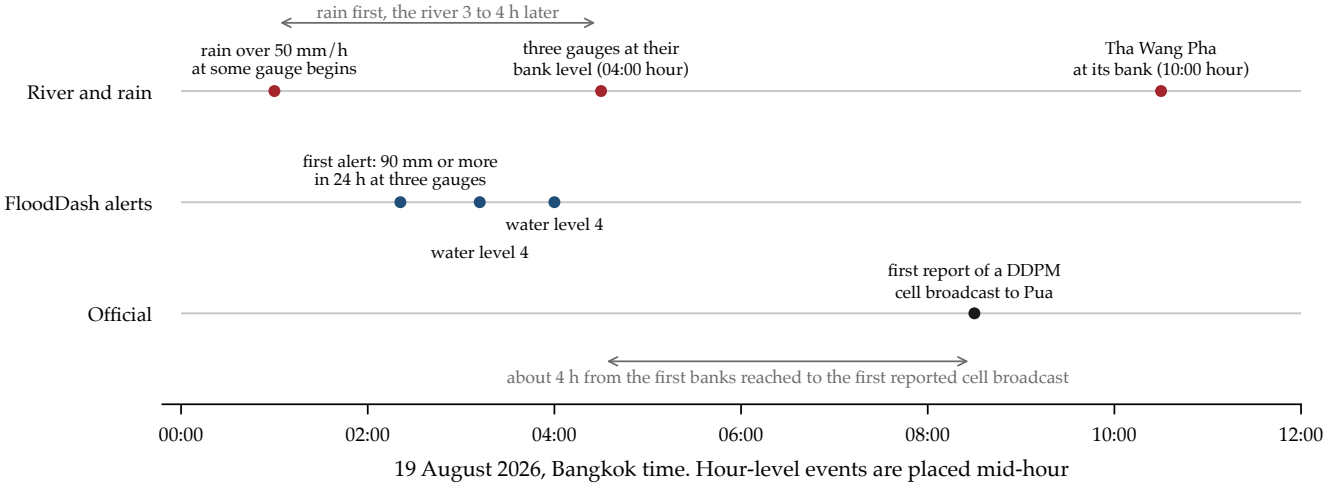


Figure 1. Nan, 19 August 2026. Rain preceded the river. FloodDash raised its first alert at 02:21 and two water-level alerts at 03:12 and 04:00. The first report of a DDPM cell broadcast is stamped 08:30. Sources: system alert table; archived hourly gauge aggregates; Thairath (Thairath, 2026). FloodDash alerts went to its own subscribers only (Section 4.8).

about 90% of damage and loss fell on the private sector. The World Bank puts roughly 70% of the damage and loss on manufacturing, after six industrial estates in Ayutthaya and Pathum Thani flooded. Climate studies have examined the causes and return periods of the rain (Gale and Saunders, 2013; Promchote et al., 2016), and supply-chain studies follow how the flooding spread through production networks (Haraguchi and Lall, 2015).

None of that is a data problem alone. But a flood this slow, spread over five months, is exactly the kind a joined-up monitoring chain should be able to describe while it is happening.

2.2 *Why warnings fail*

The warning literature is old and consistent. A warning works when it is timely, specific, from a source people trust, and given in a form they can act on (Mileti, 1999; Sorensen, 2000). Systems fail at the joins: between monitoring and forecasting, between forecasting and a decision, between a decision and a person (Basher, 2006). When the chain works, the returns are large; the monetary benefit of European flood warnings is large (Pappenberger et al., 2015). The Sendai Framework made the point a global target: increase the availability of and access to multi-hazard early warning systems by 2030 (UNDRR, 2015). A 2022 review found that countries with limited early warning coverage had about eight times the disaster mortality of countries with substantial coverage (UNDRR and WMO, 2022).

Two recent lines of work matter here. Global models now forecast extreme floods even in basins with no gauges (Nearing et al., 2024; Alfieri et al., 2013; Harrigan et al., 2020), so the science of the first link no longer waits for a national gauge network. And unprecedented floods and droughts challenge risk management that rests on past records (Kreibich et al., 2022), which is an argument for watching the present as well as the archive.

2.3 *What a bypass must not do*

"Bypassing bureaucracy" can sound like a boast, so the limits come first. Flood management is a wicked problem in the sense of Rittel and Webber (1973): no single owner, no clean boundary, no test run. Societies also manage shared resources through many overlapping centres of authority, and sometimes do it well (Ostrom, 1990). Flood defences and settlement patterns interact, and the interaction can raise exposure over time (Di Baldassarre et al., 2013). An independent warning system is therefore one more node in a crowded field, not a replacement for it. It has no power to order an evacuation and no right to be believed. It earns trust only by being open about its sources, its errors and its limits. Volunteered and independently built tools can add speed and reach in a crisis (Zook et al., 2010), and they can also add noise. We return to that in Section 4.9.

3 The system and how we measure it

3.1 *Design*

FloodDash has four layers (Figure 2). *Ingest* polls 44 public feeds on schedules from five minutes to a day and archives every reading. *Judge* turns readings into alerts with a small set of rules. *Say* sends alerts to people. *Check* watches the other three.

The judging rules are few and plain. A water-level alert fires when a ThaiWater gauge reports situation level 4 (near bank) or 5 (over bank). A precursor alert fires when a gauge is rising fast toward its bank. A heavy-rain alert fires when a cluster of gauges passes an intensity line that drops on soaked ground. A 24-hour rain alert fires on large accumulations. A relay re-sends Thai Meteorological Department (TMD) warnings, and since 2 November it forwards a TMD warning to a province's subscribers only after a gauge there has measured at least 10 mm h^{-1} (Thai Meteorological Department, 2026). A wetness engine, built on antecedent rainfall, lowers the thresholds when the ground is al-

ready full.

Five honesty rules apply to everything the system says. A score is never called a probability. A reading carries its source and its time. A missing reading is never presented as a safe one. Nothing is described as passable because a sensor reads shallow water. And no sentence states a depth at a house, because the elevation data we hold cannot support one.

3.2 Data

The system draws on 44 feeds from 30 publisher labels, several of which are one body under two names. Twenty-eight feeds come from Thai public bodies: the Hydro-Informatics Institute (HII), TMD, the Geo-Informatics and Space Technology Development Agency (GISTDA), the Bangkok Metropolitan Administration (BMA), the Pollution Control Department, the Royal Irrigation Department, the Department of Water Resources, and two municipalities. Ten come from eight international providers: Copernicus services, JAXA, NASA, NOAA, USGS, the National Hurricane Center, the Mekong River Commission and Open-Meteo (Zippenfenig, 2024; Hersbach et al., 2020). Six are neither: a traffic foundation, a news aggregator and our own derived pipelines.

Rivers are routed with HydroRIVERS (Lehner and Grill, 2013; Lehner et al., 2008), which we integrated on 4 November under its licence terms (HydroSHEDS, 2026). Satellite flood extent is read from the Copernicus Global Flood Monitoring service (Chini et al., 2022). Thai water-level and rain gauges come through the HII ThaiWater platform (Hydro-Informatics Institute, 2026).

The database holds 33.9 million raw readings and 19.6 million hourly aggregates. A retention policy prunes raw readings, so for older events we use the hourly aggregates. Ingest-run logs cover 19 September to 4 November 2026; the alert table covers 3 July to 4 November.

3.3 Measures

We report nine things, each from the system's own tables, regenerated by one script (Data availability): (i) feed reliability, as the share of successful runs; (ii) how much of the network reported recently; (iii) whether gauge metadata pass a physical sanity check; (iv) agreement between two agencies' rain gauges less than 2 km apart, over 21 days; (v) latency, as the gap between a gauge's own timestamp and the alert it triggered; (vi) forecast skill, from a daily replay of the precursor engine over the retained 40-day window, scoring a station as a hit if the watchlist flagged it and its situation level then reached 5 within 48 hours; (vii) two case reconstructions; (viii) reach, as subscriber counts; (ix) a census of the citizen-built flood tools we have studied.

3.4 What we cannot measure

Three limits shape everything that follows. First, the skill measure scores the system against the *same gauge network it reads*. It tests coherence, not accuracy against an independent flood record (Murphy, 1993). Second, the replay is not the stream anyone received: several of its signals did not exist on the nights we examine. Third, we cannot observe coordination inside agencies. We see only its symptoms at the data interface. H1 can therefore be described here, not tested.

4 Results

4.1 Fragmentation is visible at the interface

Forty-four feeds from 30 publisher labels is the first number. The feeds do not share a station identifier, so the system keys every station by publisher and publisher-assigned code. The formats vary from JSON services to RSS indexes that point at XML files. TMD's alert address says CAP but returns an RSS index; each item links to a separate Common Alerting Protocol file. BMA's canal service needs a POST with the right referer and refuses requests at random. A Bangkok drainage bulletin is a PDF whose text layer once broke the parser on one corrupt cell. GISTDA's gateway lists 42 services, every one of which needs a key. One BMA page that listed roads to avoid stopped returning its road list and has not recovered.

These are small irritations taken one at a time. Together they are the cost of coordination failure paid in engineering hours by anyone who wants to join the data. They are also a measurable barrier: each is a place where a second team will have to do the same work again.

4.2 Reliability is good, and the exceptions are instructive

Across 128,087 runs of 43 feeds between 19 September and 4 November, 97.2% succeeded (Figure 3). The median feed succeeded 98.6% of the time. Nineteen feeds passed 99%, and six fell below 90%. The worst was a GISTDA soil-moisture service that never succeeded, because we have no subscription to it. Three feeds of forecast and satellite rain (IMERG, a national rain grid, Open-Meteo) succeeded 59–75% of the time. The Bangkok drainage bulletin succeeded 87.5%.

The whole system was down three times: 8.3 hours on 20 October, 7.4 hours on 22 October and 31.5 hours from 24 October. Together that is 47.2 of 1,090.4 observed hours, an availability of 95.7%. The 31.5-hour outage began with a host panic after the computer ran out of memory, and it covered the day before Phuket flooded (Section 4.7). Counting only each feed's own silence, outside those outages, 21 feeds went quiet for more than six hours at least once and five for more than 24.

The system's own 30-minute self-audit, 51 checks each time, passed 86.1% of 122,617 checks, warned on 7.4% and failed 6.5%. Most failures are known and permanent, such as the key-gated feeds. We publish the number because a system that audits itself and shows only passes is not auditing.

4.3 The network is large, and half of it is stale

The system catalogues 10,412 stations and has at least one reading from 10,410. That sounds like coverage

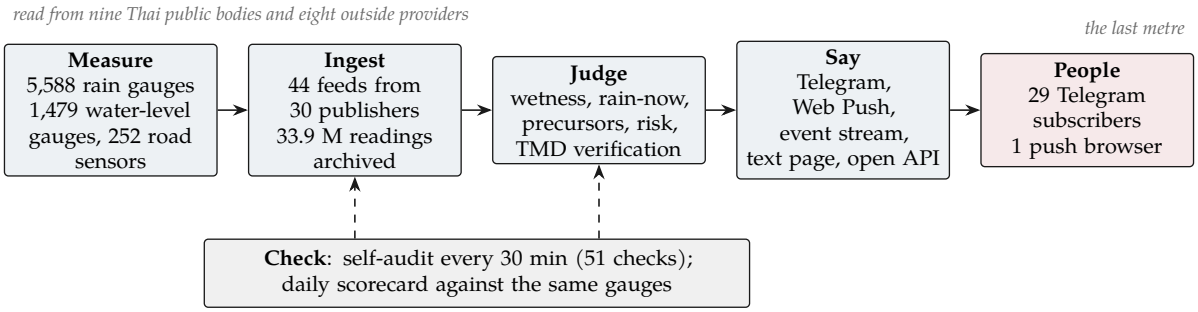


Figure 2. Architecture. The shaded blocks run unattended on one computer. The right-hand block is where Section 4.8 finds the failure.

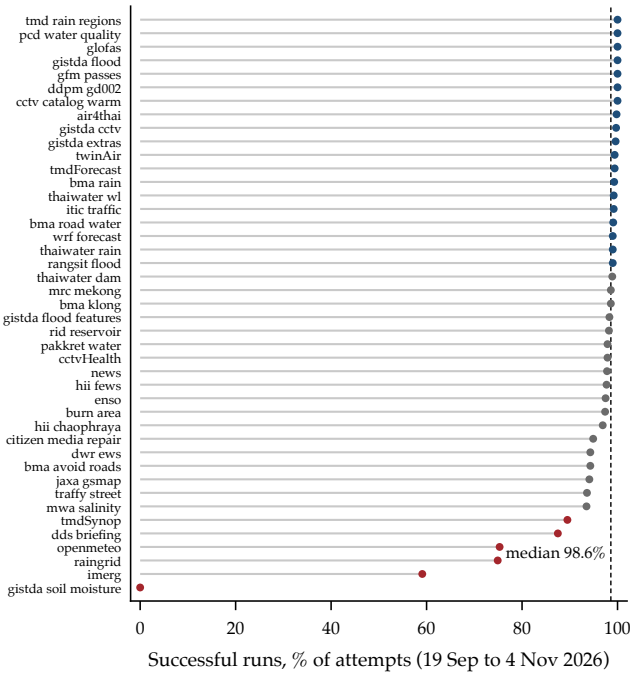


Figure 3. Reliability of 43 polled feeds, 19 September to 4 November 2026. Red: below 90%; grey: 90–99%; blue: 99% or more. Source names are the system's own identifiers.

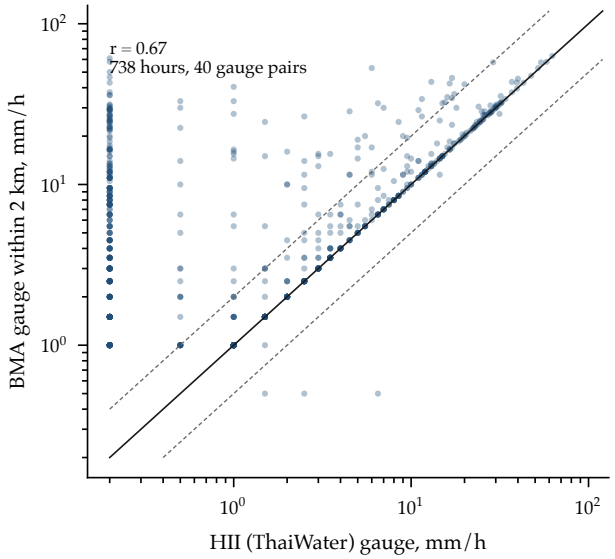


Figure 4. Two agencies, one rain. Hourly peak rain rate at 40 BMA gauges against the nearest ThaiWater gauge within 2 km, 21 days. The solid line is agreement; dashed lines are a factor of two.

until you ask who reported recently. Of 5,588 ThaiWater rain gauges, 79% reported in the last three hours. Of 2,273 early-warning stations of the Department of Water Resources, 86% did. All 126 BMA rain gauges and 128 of 129 TMD synoptic stations did.

The water-level gauges are the weak link. Of 1,479 ThaiWater river gauges, 733 (49.6%) reported in the last three hours and 770 (52.1%) in the last 24. Nearly half of the river network that a flood warning most needs was silent when we looked.

4.4 The agencies' numbers do not always agree

Every one of the 1,479 water-level gauges carries a stated bank level. We compared each with its latest water level. A bank that sits more than 30 m above the water, or more than 2 m below it, cannot be right. Thirty-six gauges (2.4%) fail the test. The error is small in share and large in consequence, because a wrong datum turns a river at its bank into one 30 metres below it.

Rain gives a sharper picture. We matched 126 BMA rain gauges to the nearest ThaiWater gauge within 2 km and compared the 40 closest pairs over 21 days (Figure 4). In the 738 hours in which either gauge read at least 1 mm h^{-1}, the two agree with a Pearson correlation of 0.67. In the 216 hours when both read at least 5 mm h^{-1}, the median ratio is 1.0 but 10.6% differ by more than a factor of two. In 136 hours one gauge read 5 mm h^{-1} or more while the other read under 1. Some of that is real: a convective cell can drop 5 mm h^{-1} on one side of 2 km and not the other. Some is timing, since each gauge reports on its own clock. Either way, a warning built on one agency's gauge alone can look very different from one built on the other's.

4.5 Speed: minutes from gauge to alert

For 162 water-level and precursor alerts we matched the alert time to the gauge's own timestamp (Table 1). The median lag is 6.7 minutes. A quarter of alerts follow within 4.6 minutes and nine in ten within 54 minutes. The lag includes the publisher's own delay, our polling interval and processing, and excludes delivery to a phone.

The system raised 13,766 alerts in 124 days, 111 a day (Table 1). That is too many to push unfiltered. Of them, 1,792 were severity 3. Telegram and Web Push send

Table 1. Alerts and latency, 3 July to 4 November 2026.

Alerts raised	13,766
water level (level 4 or 5)	5,143
flash-flood precursor	4,024
24-hour rain	3,270
heavy rain now	669
reservoir and dam	365
near-bank rush	157
outlook and TMD relay	138
Severity 3	1,792
Province band changes	1,369
at night (22:00–06:00)	37
Gauge-to-alert lag, median (n=162)	6.7 min
25th and 90th percentile	4.6, 54.4 min

alerts by province and severity, and a TMD warning is forwarded only after a gauge in the named province confirms it. In its first two days that relay forwarded 25 verified alerts.

4.6 Skill: the watchlist catches most overflows and is wrong most of the time

Over the retained 40 days (25 September to 4 November), 199 gauges served as the test set. The 48-hour watchlist flagged 302 of 358 station-days that went on to reach level 5, a recall of 84.4%. It also flagged 2,115 that did not, a precision of 12.5%. The median lead was 8.2 hours, with a quarter of hits flagged 3.4 hours ahead or less and a quarter 25.4 hours ahead or more. With the live rain signal switched off, leaving only wetness, forecast and bank height, recall was unchanged at 84.3% and the median lead 7.7 hours. That second number matters more. It is the warning that exists before the rain falls.

Seven false alarms to every hit is a poor ratio for an alert and a defensible one for a watchlist. We built the list wide on purpose: missing a flash flood costs lives, and an extra row on a screen costs a glance. It also means the list must not be pushed raw to people, which is why the push channels apply the filters described above. Recall, precision and lead are the same measure the system publishes continuously, and an earlier run on 9 September gave 74.2%, 3.8% and 19.7 hours. The numbers move as the season does.

4.7 Two nights

Nan. Figure 1 gives the timeline. Rain at some gauge passed 50 mm h^{-1} in the 01:00 hour and 80 in the next. Three river gauges reached their stated bank levels in the 04:00 hour. FloodDash raised its first alert at 02:21, for 90 mm or more in 24 hours at three gauges, and water-level alerts at 03:12 and 04:00. The first report of a DDPM cell broadcast to Pua is stamped 08:30. A gap of about four hours separates the river reaching its banks from that report. FloodDash's own alerts reached only its subscribers, and the heavy-rain cluster rule that would have fired at 01:00 did not yet exist. We built it afterwards, and replayed on the Nan data it fires at 01:00.

Our own earlier write-up of this night also said that the system raised a level-5 alert at 04:11 and moved the province to its top band at 03:13. We checked those two claims against the alert and escalation tables on 4 November and could not find either. We have removed them. We mention it because it is the failure this paper is about, in miniature: a record that said more than the data did.

Phuket. On 25 and 26 October the system was down for 31.5 hours, from 24 October 15:42 to 25 October 23:10. It restarted and raised a 24-hour rain alert at 23:12. At 01:00 the strongest of 17 Phuket gauges read 47.8 mm h^{-1}. A flash-flood precursor fired at 01:41, a water-level alert at 02:53 and the heavy-rain alert at 03:25 (Figure 5). Thai news reports that morning describe a flash flood of mud and rock through Patong in the middle of the night. The heavy-rain rule had required three gauges over 20 mm h^{-1} or one over 50 with a second heavy gauge, and one cell over steep ground met neither. We added a rule on 27 October for saturated ground: one gauge at 40 mm h^{-1} or more and a second at 10 or more. Replayed on 15 September to 27 October it adds about 1.7 province-days a day, 30% of them matched by flood news against 36% for the old rule, and it fires at 01:00, two hours and 25 minutes before the 03:25 alert. No gauge sits in Patong itself.

4.8 The last mile: 10,410 stations, 29 people

This is where H3 breaks. The system has 10,410 reporting stations, 13,766 alerts and 1,792 severity-3 alerts. It has 29 Telegram subscribers, 23 of whom chose a province, covering 12 provinces. One browser has subscribed to Web Push, which went live on 1 November. LINE, the dominant messenger in Thailand, is closed to us: its free tier allows 300 messages a month and bills per recipient, so a single broadcast to 300 followers consumes the month, and we retired it.

On 2 November TMD warned of heavy to very heavy rain in the south for 5 to 9 November and named seven provinces. We counted subscribers in those provinces. There were none. The engines were ready and nobody was listening. Figure 6 puts the numbers on one chart.

4.9 The citizen-built tools

From September to November we studied 48 other flood tools, found mostly through social media and word of mouth. Twenty-eight are volunteer projects, ten are research tools, six come from companies, one from a political party and three have no named operator. Twelve state or show that an AI coding tool was used. Seven plan routes around water, and 19 invite citizens to report it. On our own reading of their evidence practices, 23 rate high for honesty, 16 medium, three mixed and six are unrated; the rating is our judgement and is published card by card.

They are the same unmet need showing up 48 times. Each one re-ingests overlapping public feeds. Each one reaches a few people. We took ideas from 25 of them and credit each one on its card. If the official side offered

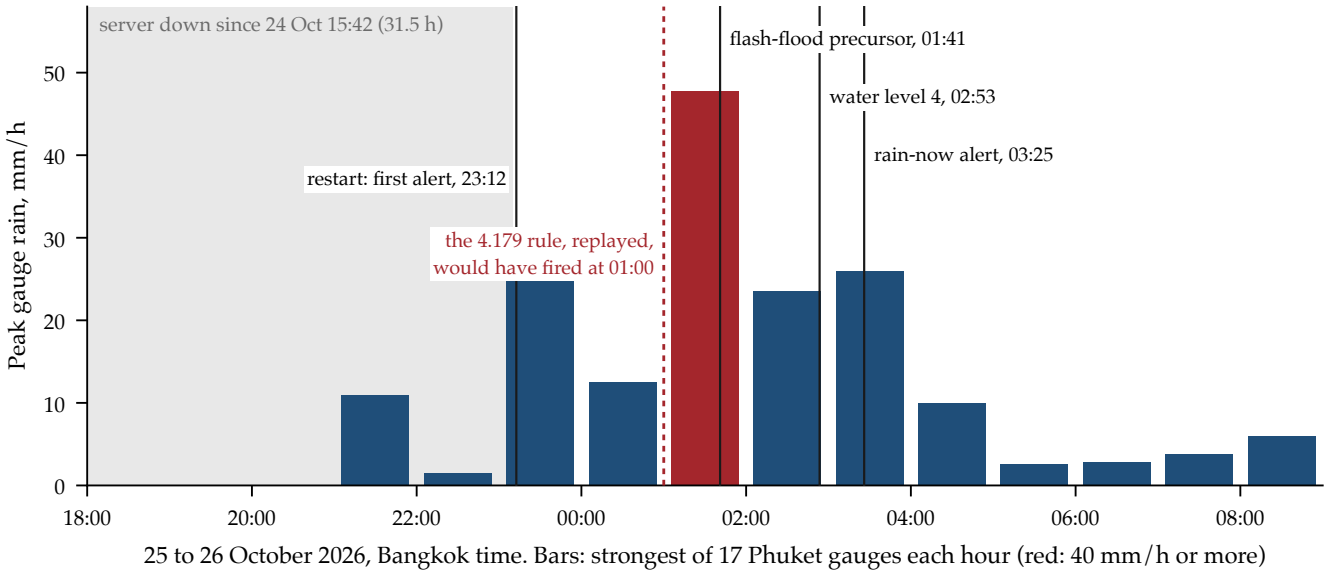


Figure 5. Phuket, 25 to 26 October 2026. The server was down for the first part of the window. Alert times from the alert table; rain from raw gauge readings.

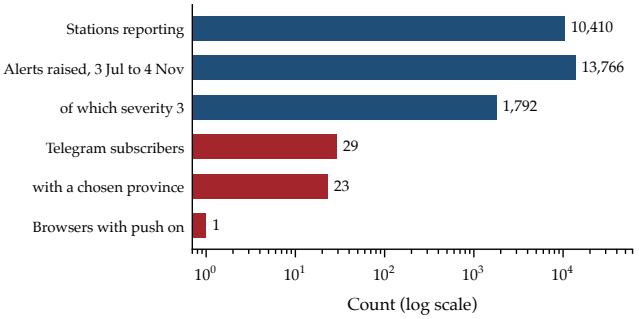


Figure 6. From measurement to people. Blue: what the system can do alone. Red: what needs people to opt in.

a stable open interface, most of that effort could go into reach.

5 Discussion

5.1 What the evidence says about the three hypotheses

H1, coordination. We can describe it and cannot prove it. At the data interface the signs are plain: 44 feeds under 30 publisher labels with no shared station identifier; gauges from two agencies 2 km apart that correlate at 0.67; 36 bank levels that cannot be right; a river network in which half the gauges had gone quiet. These are symptoms. They are what you would expect if each agency built its system for its own mandate. They would also appear, weaker, in any well-run set of agencies. We cannot see the meetings that did not happen.

H2, responsiveness. The data are fast enough. A median of 6.7 minutes from gauge to alert is not the bottleneck. Recall of 84% over 48 hours is workable. The responsiveness that failed on 19 August and 26 October was not speed but rules: the heavy-rain rule that needed breadth, and a headline that was 40% water level in a basin where water follows rain. Responsiveness is a design problem that a small team can fix in days, and did.

H3, the chain. The monitoring-to-alert half exists and runs. The alert-to-person half is nearly empty. That is the main negative result. The system is also fragile in the ordinary way of a one-computer project: 95.7% availability, with its longest outage the day before Phuket flooded.

5.2 What a bypass can and cannot bypass

Bypassing worked for three things. It needed no permission to read open data, so integration took weeks and cost nothing. It needed no committee to change a rule, so the Phuket fix shipped within two days of the event. And it needed no budget cycle to publish its own errors.

It could not bypass the thing that decides whether anyone is warned: the right to reach a phone. The DDPM's cell broadcast reaches everyone in a cell tower's footprint without anyone having opted in. Our channels reach only people who found us first. No independent system, however accurate, will reach the people a cell broadcast reaches. The honest conclusion is not that independent systems should replace the official channel. It is that they should feed it, and the official channel should be able to take their input.

The bypass is also not independent of the agencies. It lives on their data. If ThaiWater closes its API or a gauge stream goes dark, FloodDash goes blind. A bypass built on open data protects the public only as long as the data stay open.

5.3 A minimum open interface

Our results point to a short list that would do more for coordination than another committee. The list is small on purpose.

1. Every public gauge feed carries a stable identifier, a declared datum, and the time of the measurement, not only the time of publication.
2. Every public feed states its cadence, and publishes its

own staleness, so a client can tell a quiet river from a dead gauge.
3. Public warnings are published in the Common Alerting Protocol, with province or polygon codes, as TMD already does for rain.
4. The official channel accepts machine-readable *draft* alerts from independent systems that meet a published honesty standard, and labels them as unofficial. A relay of this kind would let a system like this one add its verification to an official warning without claiming to be one. We built a small version of it for TMD warnings, and it works because TMD publishes a machine-readable feed.
5. Agencies publish their own forecast skill, as a ledger, the way we do.

5.4 Limitations

The study has one system, one operator and 124 days, at the tail of one flood season. It cannot show that a coordinated agency system would perform better, or worse. The skill measure checks the system against its own gauges and replays its engine over retained data. Our two cases were chosen because they were notable, and notable cases flatter a method. The case record is not free of error either, as the correction in Section 4.7 shows. The system is not an official warning and says so on every screen. We count ourselves among the 48 tools in Section 4.9.

Finally, the research question presupposes a failure. We cited the two worst floods because they are the clearest, not because they prove the cause. A fuller test needs the agencies' side of the story.

6 Conclusion

An independent system can read a country's open water data, judge it in minutes and publish its own errors, with no permission and little money. It cannot, alone, reach the people the official channel reaches. The next step is therefore not another dashboard. It is an interface that lets the official channel take verified input from systems like this one, with names attached and limits stated. We have built our half and published how it performs. The other half belongs to the agencies, and it is mostly a matter of choosing to publish three fields and accept one kind of message.

Declarations

Data and code availability. Every number in Section 4 is regenerated by `paper/analysis.mjs`, which reads the live database and the public API and writes `paper/data/analysis.json` and CSV tables; figures are drawn by `paper/figures.py`. The scripts, derived tables and the LaTeX source are published with this preprint under CC BY 4.0. The source repository was private at the time of writing; access can be requested through the correspondence address. The HydroRIVERS derived subset is not redistributed, following its licence.

AI assistance. The data-analysis scripts, figures and a first draft of the text were prepared with Claude (Anthropic) under the second author's direction. Numbers come from the scripts, not from the model. Reference details were checked against Crossref, DataCite or the publisher's page. The authors are responsible for the content, and co-author review is pending for this version.

Competing interests. The second author built and operates the system described. The views are the authors' own and not those of their institutions. No funding was received for this work.

Author contributions. To be completed by the authors.

Acknowledgements. We thank the public bodies whose open data this work stands on. Their data use does not imply their endorsement of this paper.

References

Al Jazeera. Death toll from flooding in southern Thailand reaches at least 145. News report, 28 November 2025, 2025. URL https://www.aljazeera.com/news/2025/11/28/death-toll-from-flooding-in-southern-thailand-reaches-at-least-145.

Lorenzo Alfieri, Peter Burek, Emanuel Dutra, Blazej Krzeminski, Davide Muraro, Jutta Thielen, and Florian Pappenberger. GloFAS – global ensemble streamflow forecasting and flood early warning. *Hydrology and Earth System Sciences*, 17(3):1161–1175, 2013. doi: 10.5194/hess-17-1161-2013.

Reid Basher. Global early warning systems for natural hazards: systematic and people-centred. *Philosophical Transactions of the Royal Society A*, 364(1845):2167–2182, 2006. doi: 10.1098/rsta.2006.1819.

Marco Chini et al. An ensemble-based approach to map floods globally using Sentinel-1 data: The Global Flood Monitoring system of the Copernicus Emergency Management Service. EGU General Assembly 2022, abstract EGU22-11555, 2022.

Giuliano Di Baldassarre, Alberto Viglione, Gemma Carr, Linda Kuil, José L. Salinas, and Günter Blöschl. Socio-hydrology: conceptualising human-flood interactions. *Hydrology and Earth System Sciences*, 17(8):3295–3303, 2013. doi: 10.5194/hess-17-3295-2013.

Emma L. Gale and Mark A. Saunders. The 2011 Thailand flood: climate causes and return periods. *Weather*, 68(9):233–237, 2013. doi: 10.1002/wea.2133.

Masahiko Haraguchi and Upmanu Lall. Flood risks and impacts: A case study of Thailand's floods in 2011 and research questions for supply chain decision making. *International Journal of Disaster Risk Reduction*, 14:256–272, 2015. doi: 10.1016/j.ijdrr.2014.09.005.

Shaun Harrigan, Ervin Zsoter, Lorenzo Alfieri, Christel Prudhomme, Peter Salamon, Fredrik Wetterhall, Christopher Barnard, Hannah Cloke, and Florian Pappenberger. GloFAS-ERA5 operational global river discharge reanalysis 1979–present. *Earth System Science Data*, 12(3):2043–2060, 2020. doi: 10.5194/essd-12-2043-2020.

Hans Hersbach, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, et al. The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049, 2020. doi: 10.1002/qj.3803.

Hydro-Informatics Institute. Thaiwater: national water data platform. Web service, accessed 4 November 2026, 2026. URL https://www.thaiwater.net.

HydroSHEDS. HydroRIVERS v10 and licence agreement. Web page, accessed 4 November 2026, 2026. URL https://www.hydrosheds.org/products/hydrorivers.

Heidi Kreibich, Anne F. Van Loon, Kai Schröter, Philip J. Ward, Maurizio Mazzoleni, et al. The challenge of unprecedented floods and droughts in risk management. *Nature*, 608:80–86, 2022. doi: 10.1038/s41586-022-04917-5.

Bernhard Lehner and Günther Grill. Global river hydrography and network routing: baseline data and new approaches to study the world's large river systems. *Hydrological Processes*, 27(15):2171–2186, 2013. doi: 10.1002/hyp.9740.

Bernhard Lehner, Kristine Verdin, and Andy Jarvis. New global hydrography derived from spaceborne elevation data. *Eos, Transactions American Geophysical Union*, 89(10): 93–94, 2008. doi: 10.1029/2008EO100001.

Dennis S. Mileti. *Disasters by Design: A Reassessment of Natural Hazards in the United States*. Joseph Henry Press, Washington, DC, 1999.

Allan H. Murphy. What is a good forecast? an essay on the nature of goodness in weather forecasting. *Weather and Forecasting*, 8(2):281–293, 1993. doi: 10.1175/1520-0434(1993)008<0281:WIAGFA>2.0.CO;2.

Grey Nearing, Deborah Cohen, Vusumuzi Dube, Martin Gauch, Oren Gilon, et al. Global prediction of extreme floods in ungauged watersheds. *Nature*, 627:559–563, 2024. doi: 10.1038/s41586-024-07145-1.

Elinor Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, Cambridge, 1990. doi: 10.1017/CBO9780511807763.

Florian Pappenberger, Hannah L. Cloke, Dennis J. Parker, Fredrik Wetterhall, David S. Richardson, and Jutta Thielen. The monetary benefit of early flood warnings in Europe. *Environmental Science & Policy*, 51:278–291, 2015. doi: 10.1016/j.envsci.2015.04.016.

Parichart Promchote, S.-Y. Simon Wang, and Paul G. Johnson. The 2011 great flood in Thailand: Climate diagnostics and implications from climate change. *Journal of Climate*, 29(1): 367–379, 2016. doi: 10.1175/JCLI-D-15-0310.1.

Horst W. J. Rittel and Melvin M. Webber. Dilemmas in a general theory of planning. *Policy Sciences*, 4(2):155–169, 1973. doi: 10.1007/BF01405730.

Royal Thai Government. Water resources act, B.E. 2561 (2018). Royal Gazette; English text via UNEP Law and Environment Assistance Platform, 2018. URL https://leap.unep.org/en/countries/th/national-legislation/water-resources-act-be2561-2018.

John H. Sorensen. Hazard warning systems: review of 20 years of progress. *Natural Hazards Review*, 1(2):119–125, 2000. doi: 10.1061/(ASCE)1527-6988(2000)1:2(119).

Thai Meteorological Department. Disaster alerts feed (common alerting protocol 1.2). Web service, accessed 4 November 2026, 2026. URL https://www.tmd.go.th/api/xml/CAP.

Thairath. DDPM sends cell broadcast flood warning to Pua district, Nan, after continuous heavy rain. News report, 19 August 2026, 08:30 (Thai headline, archived in the system's news table), 2026.

The Nation. Nan floods kill 2 and affect more than 16,000 people. News report, 20 August 2026 (figures as of 4 pm on 19 August), 2026. URL https://www.nationthailand.com/news/general/40070012.

UNDRR. Sendai framework for disaster risk reduction 2015–2030. Technical report, United Nations Office for Disaster Risk Reduction, Geneva, 2015. URL https://www.undrr.org/publication/sendai-framework-disaster-risk-reduction-2015-2030.

UNDRR and WMO. Global status of multi-hazard early warning systems: Target G. Technical report, United Nations Office for Disaster Risk Reduction and World Meteorological Organization, Geneva, 2022. URL https://www.undrr.org/quick/74610. Findings summarised in the joint press release "UN warns that half the world is not prepared for disasters", 13 October 2022.

World Bank. Thai flood 2011: Rapid assessment for resilient recovery and reconstruction planning. Technical report, The World Bank, Bangkok, 2012. URL https://recovery.preventionweb.net/publication/documents-and-publications/2011-thailand-floods-rapid-assessment-resilient-recovery-and.

Patrick Zippenfenig. Open-Meteo.com weather API. Zenodo, 2024.

Matthew Zook, Mark Graham, Taylor Shelton, and Sean Gorman. Volunteered geographic information and crowdsourcing disaster relief: a case study of the Haitian earthquake. *World Medical & Health Policy*, 2(2):7–33, 2010. doi: 10.2202/1948-4682.1069.