SCIENCE FOR RESILIENCE

NIED 防災科研

INTERNATIONAL DISASTER RESILIENCE LEADERS FORUM — INCHEON 2026

Session 4 · Institutionalizing AI-Driven Localized Disaster Safety Management Systems

Operationalizing AI in Local Disaster Response

— Institutional foundations, from Wajima 2024 and Uki 2026 —

Shingo Suzuki, Ph.D.

Project Director, Research Division for Social Resilience

National Research Institute for Earth Science and Disaster Resilience (NIED), Japan

8 September 2026 · Songdo, Incheon, Republic of Korea

The central claim

AI is a capability.

The institution is the municipal operating model.

ALGORITHM

Produces a signal, a forecast,
a classification,
or a recommendation.

OPERATIONS

Defines data, workflow,
timing, escalation,
and fallback.

AUTHORITY

Owns the decision,
the explanation, the override,
and public accountability.

Wajima proved the foundations; Uki built them mid-response — with AI helping.

Two earthquakes, two routes

WAJIMA 2024 — Noto Peninsula EQ · 1 Jan

M7.6 · shaking ≈ MMI X+

13,641 sheltered at peak

186 shelters at peak

10,612 homes damaged

UKI 2026 — Kumamoto EQ · 28 Jul

M7.1 · shaking ≈ MMI X+

15,575 households

lost water — 80% of city

18,212 aid claims in 4 weeks

THE 15-YEAR ROUTE — accumulated in normal operations

2007 A painful baseline
Paper maps; no digital victim registry

2008–2023 Everyday capability
Enterprise + cloud GIS, digital maps, trained staff

Jan 2024 Capability put to work
data connected — so support connected

THE 30-DAY ROUTE — built during the response

28 Jul M7.1 strikes at 16:27
HQ stood up in 8 minutes

Day 2–9 First service on existing base
a live dashboard of every aid claim

Day 10–30 With AI, built into a capability
12 systems in 20 days — running every morning

Readiness accumulates in peacetime — or must be built mid-response.

Connecting data, connecting support

CASE 1 Safety tags → damage certificates

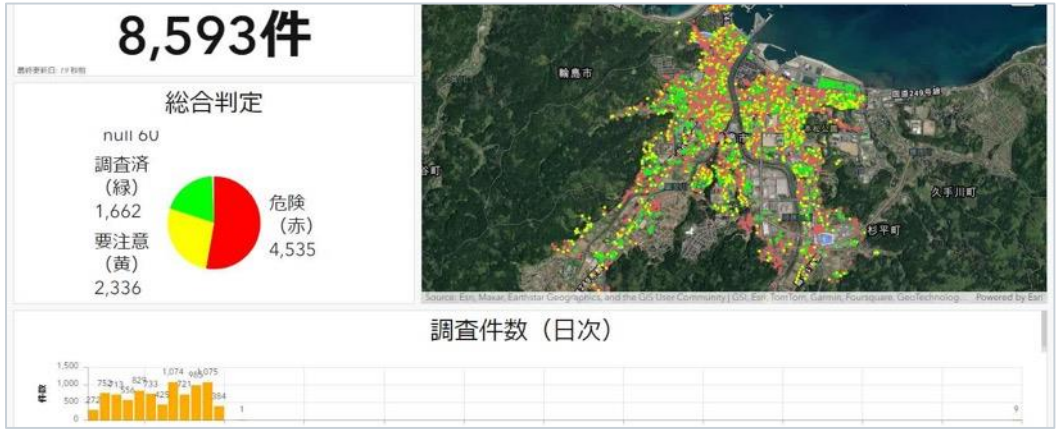

Post-quake safety tagging — 8,593 buildings

2,251 candidates → **764** certified destroyed in 2 days

≈38 team-days of field surveys saved

control: only 44% matched — human review stayed

CASE 2 Field surveys → housing decisions

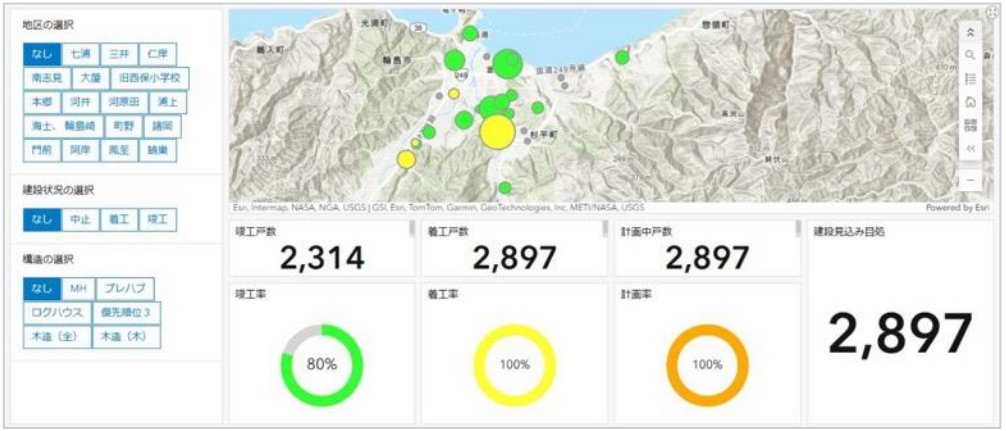

Temporary-housing construction dashboard

4,416 applications × damage results

2,860 units decided → **94%** occupied

decided before certificates were issued

The loop closed — certification results, with locations, steered household follow-up in the field

Reuse the data — but align standards, keep a shared key, keep human review.

Capability built during the response

Uki City HQ meeting — dashboard on the screen

Day 2–9 · first service

On the existing base — the aid-claims dashboard, the papers households need to unlock aid

Day 10–30 · 12 systems in 20 days

cloud-GIS data links · always-on scripts · public web apps

Running every day since

Auto-refresh with failure alerts — drop an Excel file today, see it on the morning screen

Small, fast, and embedded in the city's daily rhythm — built while responding.

Every morning, the screen sets the day

77 tasks × 65 days

the whole response structured

393 staff per day at peak

21,867 person-days in total

6:00 a.m.

refreshed daily — shown as-is
at the HQ meeting

Evidence for borrowing staff from other cities — 302/day at peak

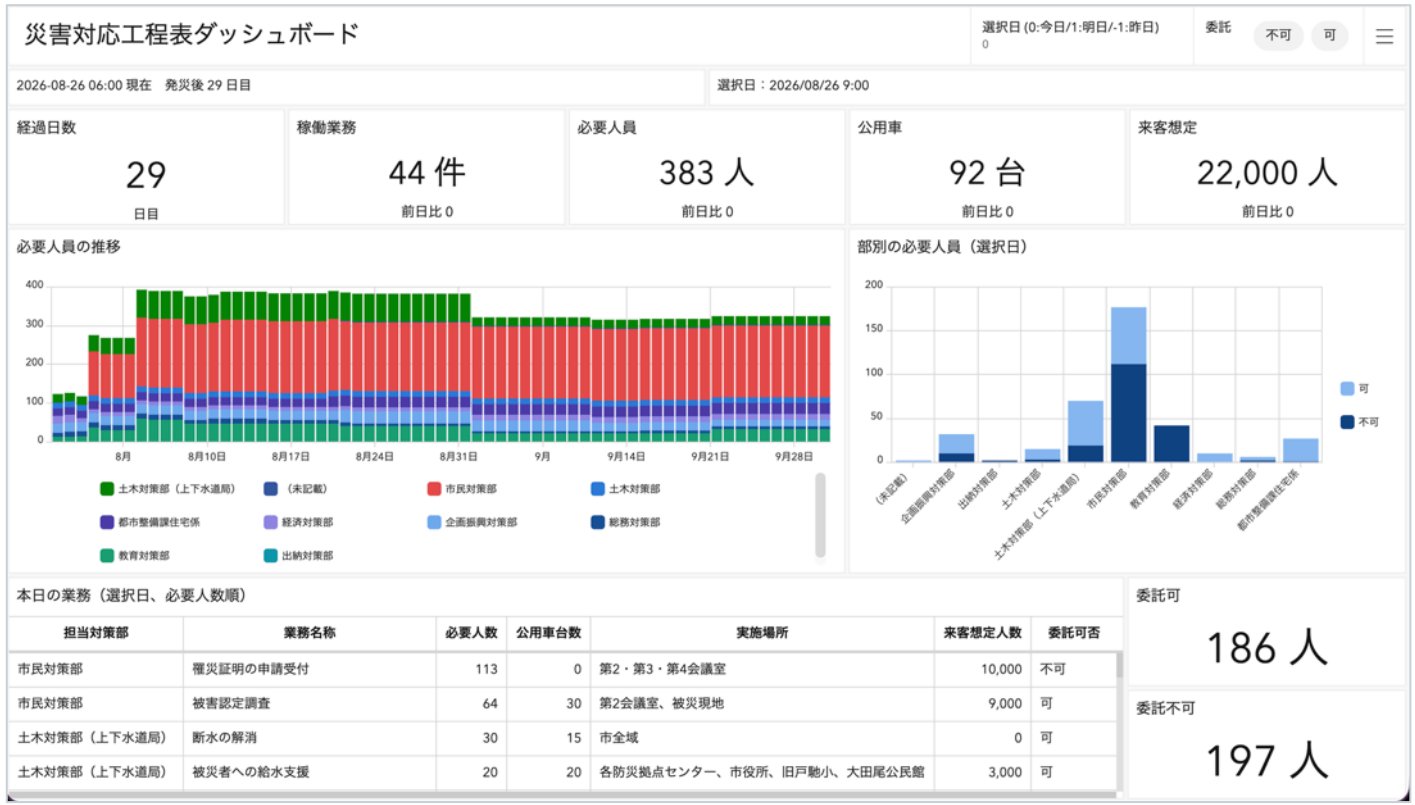

Day 29, the real screen — 383 staff needed · 44 tasks running

“How many people tomorrow?” — answered at a glance, every morning.

Forecast the whole from 16% surveyed

At 16.0% progress (1,030 of 6,428 surveyed):

3,234 cases half-destroyed or worse

90% range 3,032–3,390 — about 3× the early satellite-based damage estimate (the official grades that unlock major aid)

Why it matters

Housing, repair, and demolition are volume × unit cost × time. Resourcing starts before the surveys finish (completion due in September).

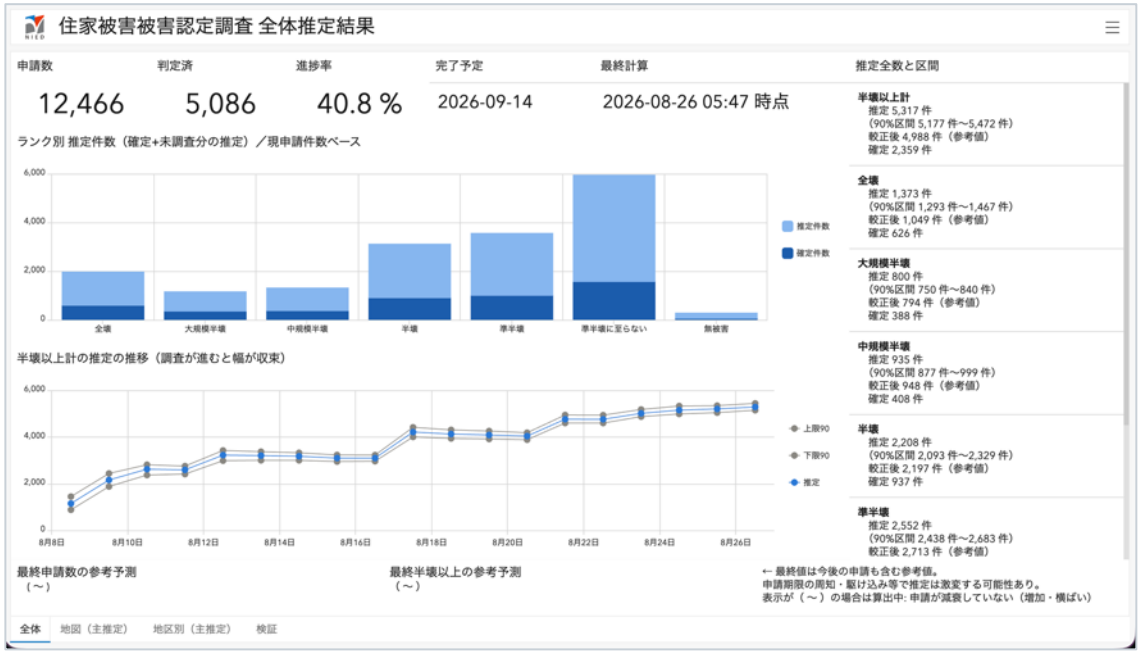

Estimation dashboard — the real screen, 26 Aug

By 26 Aug: survey caseload 12,466 · 40.8% done → 5,317 [5,177–5,472]. Recalculated at 5:00 a.m. daily, calibrated against results.

From “wait until the surveys finish” to “answer now, uncertainty shown”.

Machines flag errors; people decide

Address ↔ location

Checked against the city's 186,000-parcel address register — 11,698 OK, 124 flagged

Duplicate detection

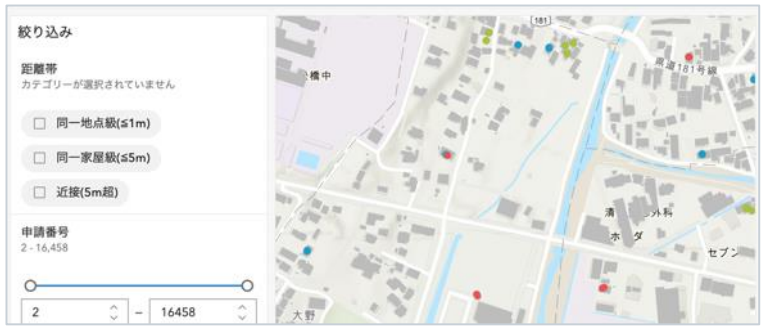

Candidate pairs within 10 m, ranked by distance — cleared the day before each survey

Photo locations

EXIF of ~60,000 photos — mix-ups and misplaced points flagged for review

Built with AI — governed by people

AI wrote the data-processing programs; the systems present evidence — they do not decide.

Assume the data is wrong somewhere — machines flag it early, people fix it.

Four roles — exercised in practice

01 NORMALIZE

Read forms, standardize addresses, link records — show confidence.

Uki: 11,000+ claims ↔ address register

02 CHECK

Flag gaps, duplicates, impossible locations, conflicting classes.

Uki: duplicate + photo-location screening

03 FORECAST

Estimate workload, demand, and recovery — uncertainty explicit.

Uki: totals estimated from 16% progress

04 SYNTHESIZE

Summarize the picture; surface bottlenecks and options.

Uki: daily auto-generated meeting page

Human authority stays explicit — approve, override, explain, log, fall back.

AI is only a capability — operations design turns it into response capacity.

Three messages — and one way of working

01

Try it

Pilot one accountable decision-support service — one decision, one owner, daily. Scale or stop.

02

Design the keys first

Common keys — person, household, case, location — and data design are peacetime decisions.

03

No disaster-only systems

A disaster-only system fails when the disaster comes — capability must live in everyday work.

One way of working — who does what

LOCAL GOVERNMENT

owns the decisions, workflows, data, and accountability

ACADEMIA & RESEARCH

co-build and calibrate on site, then hand the system over

OTHER CITIES & PARTNERS

staff the surge (up to 302/day); agree in peacetime who comes

Contact: Shingo Suzuki, Ph.D. · NIED, Japan · shingosuz@bosai.go.jp

Ask of every map, screen, and AI output: “So — which action changes?”