Seconds that Save Lives

Engineering Requirements for Earthquake Early Warning

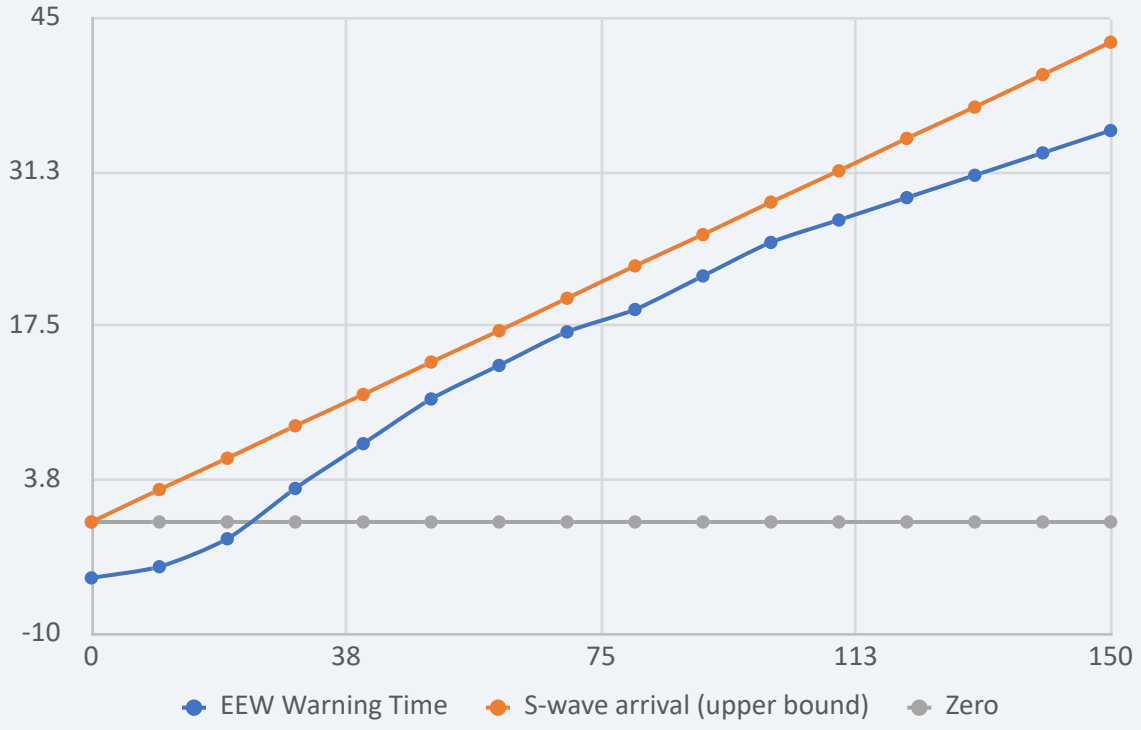
Frank Vernon

Boulder Real Time Technologies · UC San Diego

Antelope User Group Meeting · Opatija, Croatia · 3 June 2026

How much warning time is physically possible?

Minson et al. 2018 · Allen & Melgar 2019 · McGuire et al. 2025



Key distances

0–20 km
−5 to 0 s

~30 km
≈ 0 s

50 km
~10 s

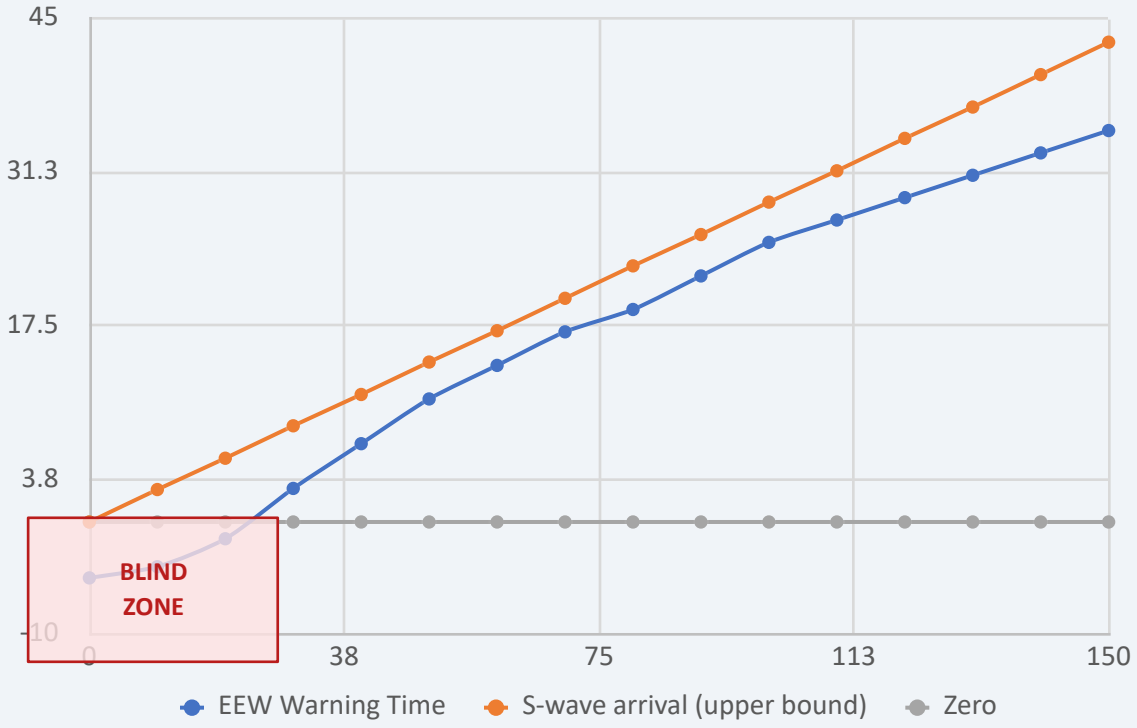
100 km
~25 s

Physical
S-wave

Every millisecond recovered from the engineering budget is a second returned to the user.

How much warning time is physically possible?

Minson et al. 2018 · Allen & Melgar 2019 · McGuire et al. 2025



Key distances

0–20 km
–5 to 0 s

~30 km
≈ 0 s

50 km
~10 s

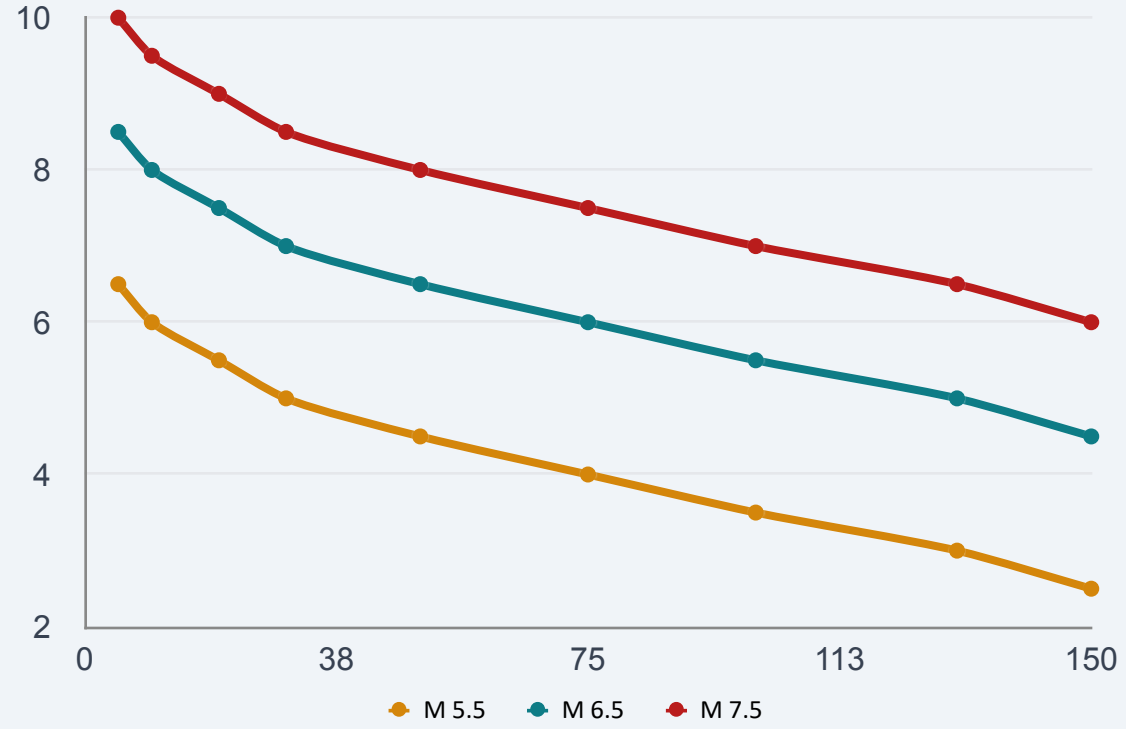
100 km
~25 s

Physical
S-wave

Every millisecond recovered from the engineering budget is a second returned to the user.

Warning time and expected shaking intensity vs. epicentral distance

Allen & Melgar 2019 (MMI attenuation) · Minson et al. 2018 (physical warning time limits)



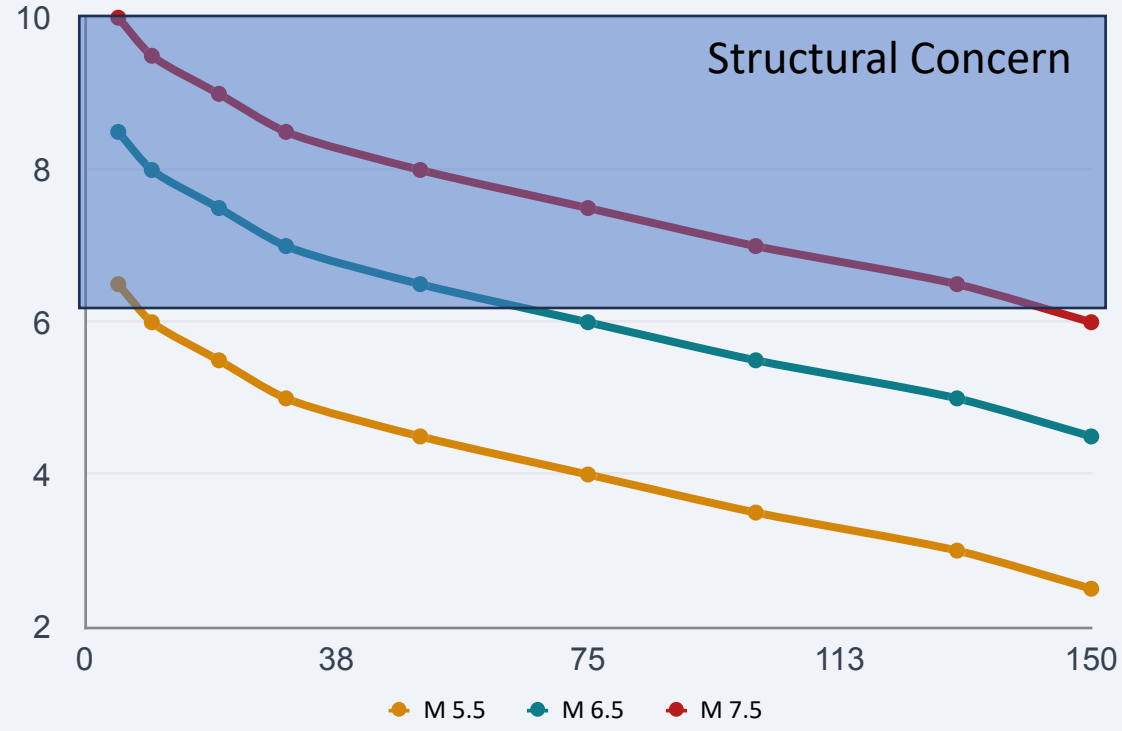
MMI at key distances

	M 5.5	M 6.5	M 7.5
10 km, -4 s	VI	VIII	IX+
30 km, 3 s	V	VII	VIII+
50 km, 11 s	IV–V	VI–VII	VIII
100 km, 25 s	III–IV	V–VI	VII
150 km, 35 s	II–III	IV–V	VI

Warning time shown at each distance. MMI scale: IV felt / VI structural concern / VIII heavy damage. Approximate — western US shallow crustal reference; local attenuation characteristics will vary.

Warning time and expected shaking intensity vs. epicentral distance

Allen & Melgar 2019 (MMI attenuation) · Minson et al. 2018 (physical warning time limits)



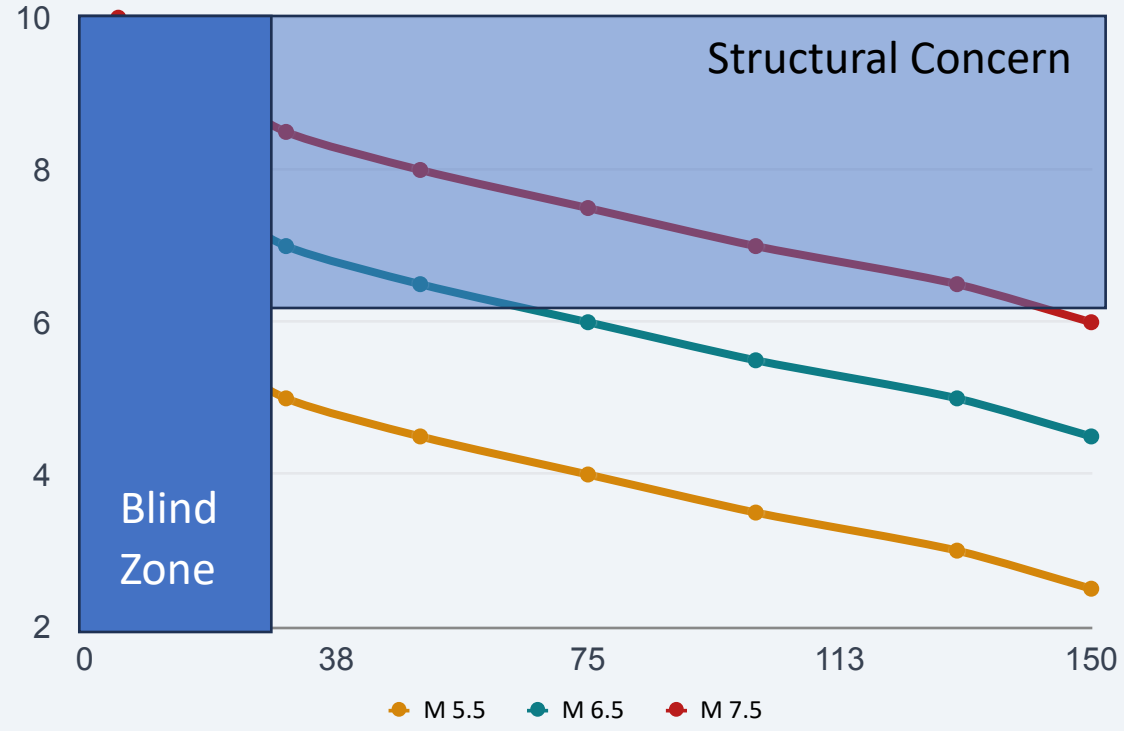
MMI at key distances

	M 5.5	M 6.5	M 7.5
10 km, -4 s	VI	VIII	IX+
30 km, 3 s	V	VII	VIII+
50 km, 11 s	IV–V	VI–VII	VIII
100 km, 25 s	III–IV	V–VI	VII
150 km, 35 s	II–III	IV–V	VI

Warning time shown at each distance. MMI scale: IV felt / VI structural concern / VIII heavy damage. Approximate — western US shallow crustal reference; local attenuation characteristics will vary.

Warning time and expected shaking intensity vs. epicentral distance

Allen & Melgar 2019 (MMI attenuation) · Minson et al. 2018 (physical warning time limits)



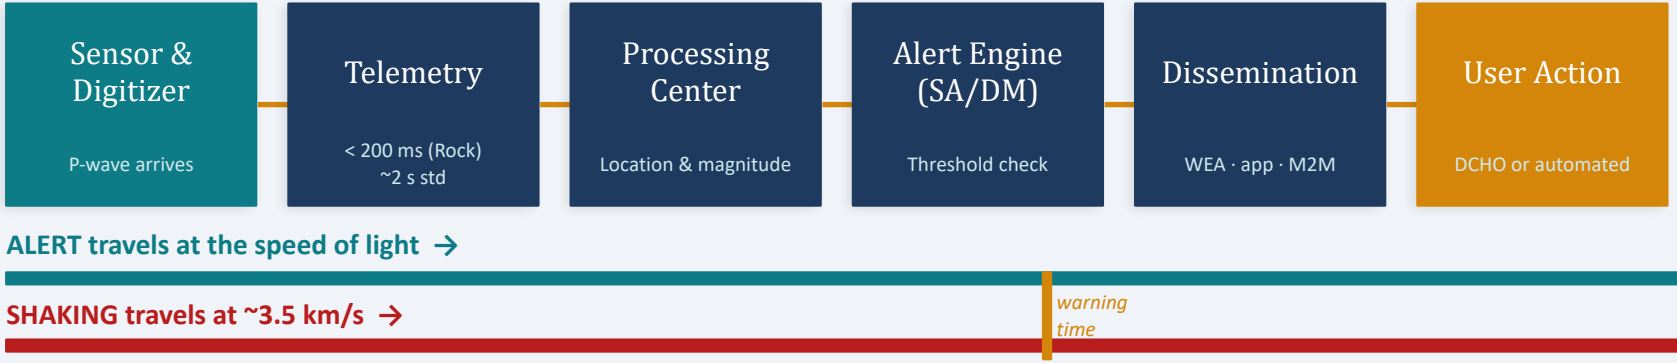
MMI at key distances

	M 5.5	M 6.5	M 7.5
10 km, -4 s	VI	VIII	IX+
30 km, 3 s	V	VII	VIII+
50 km, 11 s	IV–V	VI–VII	VIII
100 km, 25 s	III–IV	V–VI	VII
150 km, 35 s	II–III	IV–V	VI

Warning time shown at each distance. MMI scale: IV felt / VI structural concern / VIII heavy damage. Approximate — western US shallow crustal reference; local attenuation characteristics will vary.

How EEW works: the race against seismic waves

Allen & Melgar 2019



P-waves outrun S-waves

P: ~6 km/s · S: ~3.5 km/s
Detect on P, warn before S arrives

Alerts travel at light speed

Electronic signals outpace seismic waves by three orders of magnitude

Warning time = gap exploited

Alert must arrive before S-wave.
Every millisecond of latency costs warning time

The blind zone: EEW's fundamental near-field constraint

Allen & Melgar 2019 · Minson et al. 2018 · McGuire et al. 2025

Definition

The area around an epicenter where seismic waves outrace the alert — regardless of system speed.

What drives its size?

Hypocenter depth · Network station density · Minimum station count (≥4) · Algorithm processing time

Quantified constraint

8-second processing time → blind zone radius ~25–30 km for a typical shallow crustal event

Who is inside it?

The vast majority of MMI ≥8 sites for M6–7 shallow crustal events lie in the late-alert zone (Meier 2020)

How to shrink it?

Denser station networks · Lower latency hardware (Rock series) · PLUM — the only algorithm that can warn inside

MMI distances by magnitude

Source: Allen & Melgar 2019, Table 1

Mag	MMI V dist	MMI VII dist	S-wave time
M5	10–15 km	—	~5 s
M6	30–60 km	5–15 km	~15 s
M7	80–200 km	20–80 km	~40 s
M8	200–500 km	100–300 km	~2 min

What warning seconds enable: three application classes

Allen & Melgar 2019 · Minson et al. 2021 · Cochran & Husker 2019

01 Personal Safety

Drop · Cover · Hold On

5 – 15 s needed

More than 50% of earthquake injuries in Northridge and Loma Prieta were from falls and non-structural hazards — preventable with DCHO. Even a 5-second warning can halve injuries.

02 Institutional Action

Elevators · Surgeries · Gas lines

5 – 30 s needed

Automated actions with low cost-of-false-alert: elevators to nearest floor, high-speed rail braking, pipeline isolation, surgical pause. Japan has used EEW for rail since the 1960s.

03 Machine-to-Machine

Industrial systems · Infrastructure

< 1 s delivery required

Internet M2M delivers alerts in < 0.5 s. Enables protective actions for systems that cannot wait for human response. Framework: Minson et al. 2021 (Northern California rail case study).

Alert delivery technology is the binding constraint: WEA (4–30 s) · WiFi apps (< 1 s) · Machine-to-machine internet (< 0.5 s)

BLOCK 1 OF 4

Data Layer

Latency · Completeness · Telemetry · Network geometry

How data travels from sensor to processing center.
Every millisecond spent here is warning time lost at the other end.

Requirement 1: Station-to-center latency

McGuire et al. 2025 · BRTT / Kinemetrics Rock series field data

Digitizer design and packet strategy directly control how much warning time reaches the user.

Chain segment	Standard 1-second packets	Kinemetrics Rock series 100 ms packets + causal filters
Sensor digitization & buffering	~1,000 ms	~100 ms
Network telemetry propagation	~200 ms	~50 ms
Data center ingestion	~200 ms	~50 ms
Station → center total	1.4 – 2.0 s	< 200 ms
Algorithm processing (from origin)	~4 – 8 s	~4 – 8 s
Alert publication + WiFi delivery	0.5 – 1 s	0.5 – 1 s
WEA cellular delivery	4 – 30 s	4 – 30 s
Total origin → user (WiFi path)	~6 – 12 s typical	~5 – 10 s

The Rock series recovers 1–2 s of warning time vs. standard 1-second packet architectures — decisive at near-field distances.

Requirement 1: Station-to-center latency

McGuire et al. 2025 · BRTT / Kinemetrics Rock series field data

Digitizer design and packet strategy directly control how much warning time reaches the user.

Chain segment	Standard 1-second packets	Kinemetrics Rock series 100 ms packets + causal filters
Sensor digitization & buffering	~1,000 ms	~100 ms
Network telemetry propagation	~200 ms	~50 ms
Data center ingestion	~200 ms	~50 ms
Station → center total	1.4 – 2.0 s	< 200 ms
Algorithm processing (from origin)	~4 – 8 s	~4 – 8 s
Alert publication + WiFi delivery	0.5 – 1 s	0.5 – 1 s
WEA cellular delivery	4 – 30 s	4 – 30 s
Total origin → user (WiFi path)	~6 – 12 s typical	~5 – 10 s

The Rock series recovers 1–2 s of warning time vs. standard 1-second packet architectures — decisive at near-field distances.

Requirements 2 & 3: Gapless streams and redundant telemetry

Allen & Melgar 2019 · Lux et al. 2026

REQ 2 — Gapless, in-order data streams

Why gaps are fatal

A gap during rupture is indistinguishable from 'no shaking.' The processing center cannot tell whether the earthquake stopped or telemetry failed.

In-order delivery required

Out-of-order packets corrupt phase arrival timing — the foundation of all location and magnitude estimates.

Continuity benchmark

USArray achieved 99.7% data return over 10+ years: the operational target for any EEW backbone.

Consequence of failure

A single missing station during a M7+ event can expand the blind zone radius by 5–15 km.

REQ 3 — Redundant, diverse telemetry paths

Each path has a distinct failure mode. An EEW backbone must survive simultaneous loss of at least one path during the event it is supposed to detect.

FIBER

Wired / fiber

Cut by fault offset or infrastructure damage — fails when needed most

RADIO

Licensed wireless

Dedicated wireless is generally the most reliable path · Repeater chain may be vulnerable to physical damage and power outages

CELLULAR

Cellular

Overload and congestion during events; tower power loss

SATELLITE

Satellite (VSAT)

500–700 ms latency; weather-dependent; cost limits density

DAS

Distributed Acoustic Sensing (offshore)

Emerging: turns submarine fiber into seismic arrays up to 200 km offshore (MBARI SeaFOAM / ShakeAlert integration)

⚠ Cape Mendocino M7.0 (Dec 2024): 11 s information blackout — rupture ~50 km offshore, no land stations triggered until rupture was nearly complete. First alert: 15.1 s after OT.

Requirement 4: Station density and geometry

Lux et al. 2024 · Allen & Melgar 2019

ShakeAlert / Given et al. 2018 design targets:

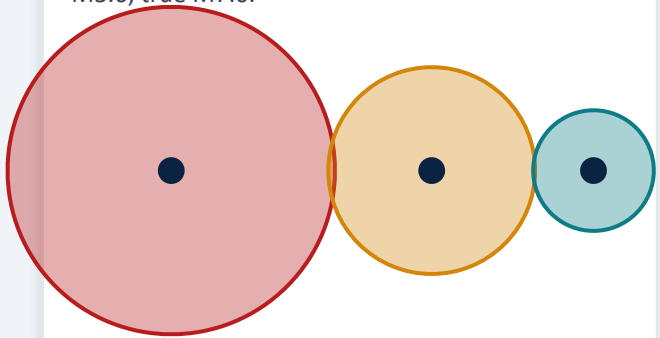
Urban / high seismicity Bay Area, Los Angeles	10 km
Seismic source regions Active fault zone areas	20 km
Lower seismic risk areas Background coverage	40 km

Algorithm detection minimums (Lux et al. 2024):

EPIC: ≥4 triggering stations · ≥40% of stations within farthest-trigger radius
FinDer: 4 neighboring stations within 50 km · PGA > 2 cm/s²

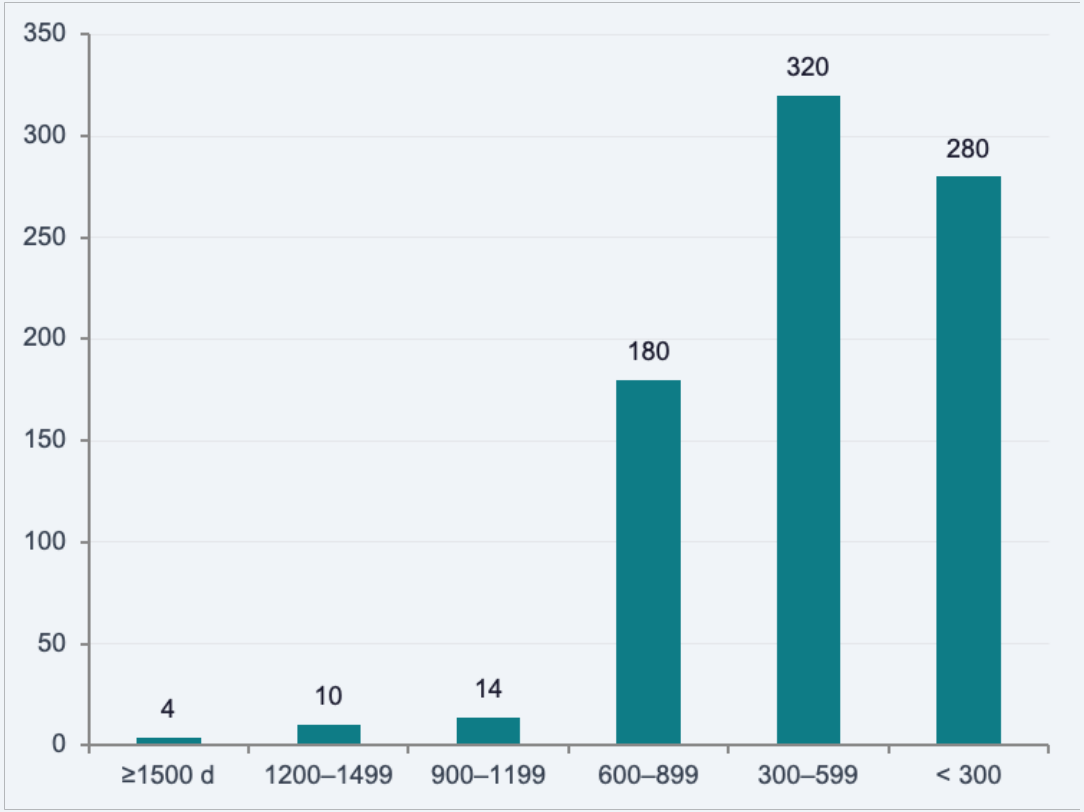
Station density shrinks the blind zone

Azimuthal coverage matters as much as density. Coastal or edge-of-network events produce one-sided station geometry — EPIC is pulled toward the network, FinDer is pushed seaward. Cape Mendocino M7.0 (Dec 2024): initial magnitude M5.6, true M7.0.



The data continuity benchmark: USArray Transportable Array

Vernon et al. · Reframed as EEW backbone target



99.7%

Data return over 10+ years

1,827 d

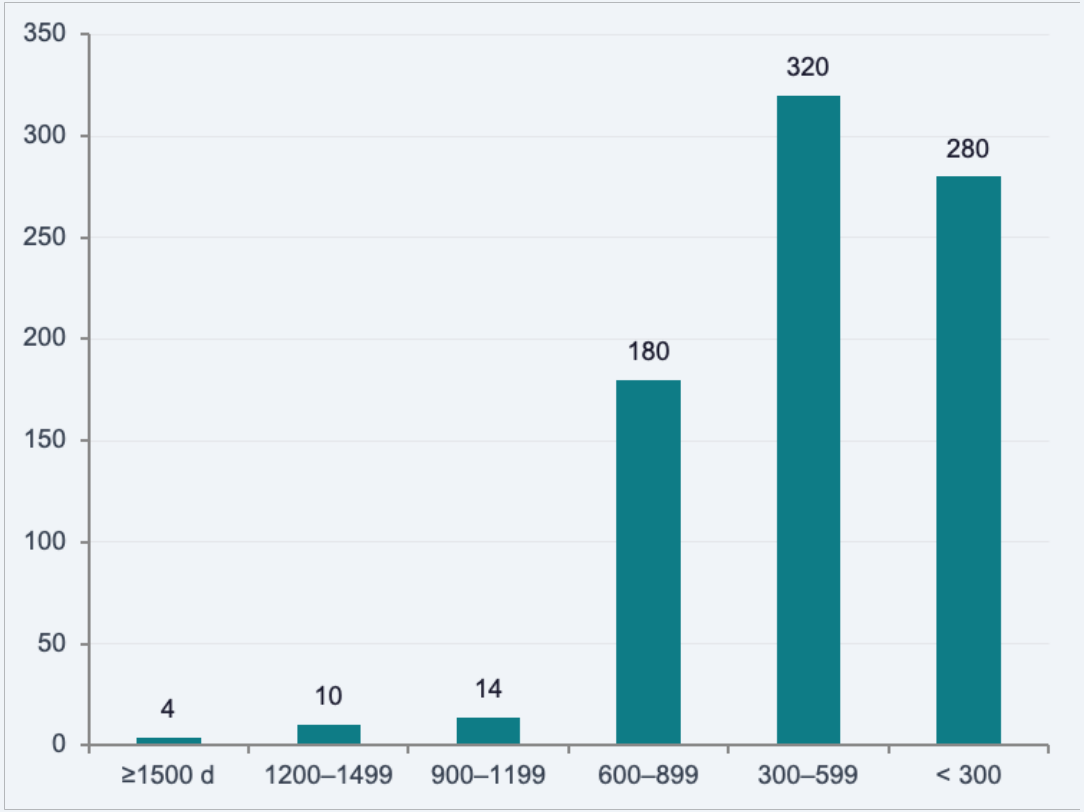
Longest continuous time series (1 stn)

416 d

Median continuous segment per station

The data continuity benchmark: USArray Transportable Array

Vernon et al. · Reframed as EEW backbone target



99.7%

Data return over 10+ years

1,827 d

Longest continuous time series (1 stn)

EEW backbone target: ~99.7% data availability should be the minimum acceptable design goal for any operational earthquake early warning network.

BLOCK 2 OF 4

Processing Layer

Algorithm diversity · EPIC · FinDer · PLUM / APPLES · GFAST-PGD

No single algorithm covers the full magnitude range — or the blind zone.
Diversity is a system requirement, not an optional upgrade.

Why algorithmic diversity is a system requirement

Allen & Melgar 2019 · Meier et al. 2020
Each algorithm class has strengths and hard limits that no single approach can bridge.

POINT SOURCE (EPIC)

Speed

Fastest first alert: 4–8 s

Strength

M4.5–M7.5 · Initial detection
Epicenter + magnitude in seconds

Limitation

Saturates at ~M7
Blind zone: cannot warn

Best applied

Small-to-moderate events
First seconds after origin

FINITE FAULT (FinDer)

Speed

Tracks growing rupture

Strength

M2–M9 · Up to 1,362 km
Line source: fault extent + strike

Limitation

Slower onset · Offshore gaps
Blind zone: cannot warn

Best applied

Large events M6.5+
Fault length characterization

GROUND MOTION (PLUM / APPLES)

Speed

Alerts from observed shaking

Strength

No source parameters needed
Immune to magnitude saturation

Limitation

Dense network required
Max lead time ~10 s

Best applied

Near-field · Inside blind zone
Complex sequences & offshore

Why algorithmic diversity is a system requirement

Allen & Melgar 2019 · Meier et al. 2020
Each algorithm class has strengths and hard limits that no single approach can bridge.

POINT SOURCE (EPIC)

Speed

Fastest first alert: 4–8 s

Strength

M4.5–M7.5 · Initial detection
Epicenter + magnitude in seconds

Limitation

Saturates at ~M7
Blind zone: cannot warn

Best applied

Small-to-moderate events
First seconds after origin

FINITE FAULT (FinDer)

Speed

Tracks growing rupture

Strength

M2–M9 · Up to 1,362 km
Line source: fault extent + strike

Limitation

Slower onset · Offshore gaps
Blind zone: cannot warn

Best applied

Large events M6.5+
Fault length characterization

GROUND MOTION (PLUM / APPLES)

Speed

Alerts from observed shaking

Strength

No source parameters needed
Immune to magnitude saturation

Limitation

Dense network required
Max lead time ~10 s

Best applied

★ **Only PLUM / APPLES can warn inside the blind zone**

Algorithm comparison: EPIC · FinDer · PLUM · GFAST-PGD

McGuire et al. 2025 · Saunders et al. 2024 · Allen & Melgar 2019 · Meier et al. 2020

	EPIC	FinDer	PLUM / APPLES	GFAST-PGD
Data input	First 4–5 s P-wave displacement	Peak accel, full event duration	Observed ground motion at nearest stations	GNSS peak displacement
Source model	Point source	Line source (length + strike)	None — bypasses source entirely	Magnitude only (seismic epicenter needed)
Magnitude range	M4.5–M7.5 (saturates ~M7)	M2–M9 up to 1,362 km	N/A — shaking-based	M7+ (triggered at SA ≥ M6)
First alert	~4–8 s after origin 4-station minimum	Can alert alone M>5.5 ~5–10 s after origin	When observed shaking exceeds threshold	Triggered by seismic algorithms
Blind zone	❌ Cannot warn inside	❌ Cannot warn inside	✅ Can warn INSIDE blind zone	❌ Cannot warn inside
Primary limit	Magnitude saturation Point source only	Sparse coverage Offshore events	Dense network required Max lead time ~10 s	No independent detection: requires seismic trigger — offshore events expose this

Quantified warning time performance

McGuire et al. 2025 · Meier et al. 2020

M7 shallow crustal — ShakeAlert V3 offline tests (McGuire 2025)

~90%	of MMI 6 sites receive 10–40 s of warning
~75%	of MMI 7 sites receive 10–40 s of warning
30 km	epicentral distance where positive warning begins
8–10 s	minimum for Drop-Cover-Hold-On (DCHO)

219 Japan earthquakes Mw4.5–9.1 — PLUM/EPIC/FinDer (Meier 2020)

40–60%	of strong-to-extreme shaking sites receive > 5 s warning
~50%	of MMI 5–6 sites (shallow crustal) receive > 5 s — all algorithms
< 20%	of MMI ≥8 sites (shallow crustal) can be alerted before shaking
PLUM	highest classification accuracy across most shaking thresholds

BLOCK 3 OF 4

Dissemination Layer

Alert delivery · Threshold design · End-user requirements · Parallel shaking

The last mile of the signal chain.
Where engineering meets human behavior — and vice versa.

Alert delivery: different users, different latency budgets

McGuire et al. 2025 · Cochran & Husker 2019

Wireless Emergency Alert (WEA)

Reaches

All WEA-capable cell phones

Alert threshold

M 5.0+ · MMI 4+

Delivery latency

4 – 30 s

Action type

Personal safety
Drop-Cover-Hold-On

Key consideration

Reaches the broadest public. Slower delivery creates tension with near-field warning times.

Cell phone apps (WiFi delivery)

Reaches

App users on WiFi / internet

Alert threshold

M 4.5+ · MMI 3+

Delivery latency

< 1 s

Action type

Personal safety
Elevators · Sirens

Key consideration

Fast enough for most protective actions. Requires user to have installed and enabled the app.

Machine-to-Machine (M2M internet)

Reaches

Automated infrastructure systems

Alert threshold

M 4.0+ · MMI 3+

Delivery latency

< 0.5 s

Action type

Train braking · Pipelines
Grid switching · Elevators

Key consideration

Internet delivery enables pre-shaking automated action. Milliseconds matter at the hardware layer.

Alert threshold is an engineering requirement, not just a policy choice

Cochran & Husker 2019 · Minson et al. 2022 · Saunders et al. 2022

The threshold choice drives backward into network design, algorithm tuning, and warning time available to the user.

TOO LOW (MMI 2–3)

Alerts for unfelt shaking → alert fatigue → users disable the system → EEW fails when it matters most.

Alert fatigue → system abandoned

OPTIMAL BALANCE

Alert when expected shaking is felt or damaging — calibrated to delivery speed, user tolerance, and warning time available.

Engineering design target

TOO HIGH (MMI 6+)

Wait for clearly damaging shaking estimate → warning time shrinks to near-zero or goes negative → alert arrives with shaking.

Zero actionable warning time

Ridgecrest 2019 M7.1: LA not alerted (below MMI 4 threshold) → public pressure → threshold lowered to MMI 3 → alerts now issued for shaking that users cannot feel. Social calibration is iterative. (Cochran & Husker 2019)

Dual-use: alerts go out while shaking observations locate damage

Bighorn structural health monitoring technology — BRTT

The same EEW backbone infrastructure simultaneously streams real-time ground motion observations for damage assessment — at no additional sensor cost.

How it works: Bighorn real-time spectral monitoring

Now-casting of wavefield spectral content

Response spectra computed continuously in real time at every network station as the event unfolds.

Continuous exceedance monitoring

Real-time comparison against engineering design criteria — alarms fire automatically when thresholds are exceeded.

No earthquake location required

Activates the moment shaking is detected at any station, before source characterization completes.

Parallel to the EEW alert chain

The same data stream that feeds EPIC/FinDer/PLUM also populates the real-time shaking observation grid.

What it delivers: real-time damage localization

Immediate post-event situational awareness

While alerts go out, the system already knows which locations experienced shaking above engineering thresholds.

First responder routing

Shaking exceedance maps available within seconds — before any ShakeMap products or field reports are available.

Quantitative engineering output

Response spectra delivered in engineering units — directly actionable for facility managers, not just seismologists.

Infrastructure operator integration

Direct interface to building control systems, emergency operations centers, and utility operators via MQTT/JSON.

BLOCK 4 OF 4

Implementation

ASPEN platform · ShakeAlert integration · Measured end-to-end latency

How the engineering requirements translate into operational systems at BRTT.

ASPEN: each engineering requirement met in production

BRTT · Kinemetrics · USGS ShakeAlert integration

✓ Sub-200 ms station-to-center latency	Kinemetrics Rock series · 100 ms packets · causal anti-alias filters · low-latency IP telemetry
✓ Gapless, in-order streaming	ASPEN ORB ringserver: continuous, gap-detected, time-ordered data delivery to all algorithm subscribers
✓ Redundant telemetry paths	ASPEN supports wired, wireless, cellular, and satellite simultaneously with automatic failover routing
✓ ≥4-station detection, correct spacing	ASPEN ingests network stations; supports ShakeAlert 10/20/40 km spacing design targets
✓ ≥99.7% data continuity	USArray operated at 99.7% over 10+ years using ASPEN architecture with SOH monitoring and remote calibration
✓ Algorithm diversity (EPIC+FinDer+PLUM)	ShakeAlert feeds ASPEN ORB; ASPEN delivers EEW alerts to buildings and infrastructure via MQTT / JSON
✓ < 1.5 s ShakeAlert → user delivery	Measured: ShakeAlert publish → ORB → OasisPlus → AWS → end user: < 1.5 s total; ASPEN chain: 400 ms
✓ Parallel real-time shaking obs.	Bighorn module computes continuous response spectra in parallel with the EEW alert chain

Measured end-to-end latency: ShakeAlert → ASPEN → user

BRTT measured data — averages over hundreds of production tests

Event / Stage	Average Delta	Average Total Elapsed
ShakeAlert publishes EEW message to USGS servers	reference	0 ms
ASPEN (Antelope) receives message from USGS	169 ms	169 ms
ASPEN posts to ORB; creates building-specific JSON packages	0.25 ms	169 ms
OasisPlus pulls ShakeAlert message from ORB	154 ms	324 ms
OasisPlus creates user alerts; pushes to AWS services	76.6 ms	400.6 ms
End user receives EEW alert via delivery channel	< 1 s additional	< 1.5 s total

Total ASPEN chain: ShakeAlert publication → end user < 1.5 seconds · Results are averages over hundreds of production tests

Engineering requirements: end-to-end summary

Layer	Requirement	What it enables	Consequence of failure
DATA	Station-to-center latency < 2 s	Preserves warning time budget	Warning time shrinks or disappears
DATA	Gapless, in-order streams	Accurate phase timing and location	False events; missed detections
DATA	Redundant telemetry (≥4 path types)	Resilience to infrastructure damage	Data lost exactly during events
DATA	≥4 stations, correct spacing	Rapid detection; small blind zone	Larger blind zone; edge-effect false alerts
DATA	≥99.7% data continuity	Algorithm reliability and training	Reduced detection probability
PROC	Algorithm diversity (EPIC+FinDer+PLUM)	Coverage across all magnitudes and distances	Blind zone coverage failure or M7+ saturation
DISS	Alert delivery < 1.5 s (WiFi / M2M)	Actionable time for DCHO and automated systems	Alert arrives with or after shaking
DISS	Calibrated alert threshold	User trust + maximum warning reach	Alert fatigue or zero warning time

The seconds are there.

The engineering is the work.

1. **Data layer**
   Sub-200 ms latency · 99.7% continuity · four redundant paths · correct station spacing
2. **Processing layer**
   EPIC for speed · FinDer for large ruptures · PLUM / APPLES for the blind zone — all three required
3. **Dissemination layer**
   < 1.5 s WiFi delivery · calibrated threshold · parallel shaking observations

Frank Vernon · BRTT / UCSD · vernon@brtt.com