



EU AI Act & GDPR Compliance Overview

Updated for legal status checked on 2 August 2026

Non-binding summary for evaluation, pilot scoping and procurement discussions

Version	Owner	Classification
v1.4	ELNAV.AI	External evaluation version

Key updates since v1.3 • Confirms Regulation (EU) 2026/1744 (Digital Omnibus on AI) is in force. • Updates the timeline: most remaining rules apply from 2 August 2026; Annex III from 2 December 2027; Annex I from 2 August 2028. • Adds the amended safety-component definition and removes reliance on the advisory label alone. • Updates the status of the draft high-risk classification guidelines after the July 2026 consultation. • Tightens retention and off-vessel data wording for external use.	How this version should be used • Suitable for external evaluation, pilot scoping, procurement and partner review. • Use with a recorded deployment-specific classification assessment for each pilot or material configuration change. • High-level enough for sharing without exposing customer-specific legal positions or technical-file detail. • Customer-specific DPIA, DPA, SOP, lawful-basis and deployment classification work remains separate.
---	--

Contact: info@elnav.ai | +385 91 443 3348

1. System boundary and intended use

Aware Mate is an on-board, human-in-the-loop active watchkeeping support system. It is advisory only and is intended to help the bridge team maintain active watchkeeping before risk escalates. The controlled product framing is: instrument the watch, not the officer.

What the system does • Uses visual and behavioural indicators without identity recognition, such as eye-closure metrics, gaze stability or direction, head pose, posture, abnormal inactivity and watch-engagement signals. • Provides proportionate local prompts for human review and action on the bridge. • Can operate offline on the vessel edge unit and can support local alarm outputs and agreed dry-contact integrations. • By default, standard off-vessel outputs are limited to aggregated or de-identified event metadata and review metrics; indirect-identification risk remains deployment-specific. • Supports privacy-bounded operational learning, pilot close-out and SOP review by authorised humans.		What the system does not do • It does not automate navigation, COLREG decisions, steering, autopilot, ECDIS or radar decision-making. • It does not perform identity recognition, biometric identification or biometric categorisation. • It does not infer emotions, stress, mood, personality or intent. • It is not designed to create personnel profiles, officer rankings, person-level risk scores or HR assessments. • It does not make automated fitness-for-duty, disciplinary, pay, promotion or dismissal decisions. • It does not require ship-to-shore raw video for the core function.	
Default processing	On-board edge processing; no cloud video is required for the core function.		
Default retention	Derived metrics and short event logs; retention is configurable under operator policy.		
Optional Critical buffer	If enabled under operator policy, a short encrypted on-device buffer may be kept for Critical alerts only, subject to documented retention, access-control and lawful-preservation rules.		
Core governance boundary	Human oversight remains central: alerts are advisory only, with mute, disable, override and review handled under the operator-approved SOP.		
Intended-purpose control	All instructions, pilot documents, sales materials and integration notes should use the same intended-purpose and excluded-use wording. Re-purposing the system may change its AI Act classification.		

2. EU AI Act status as of 2 August 2026

Current public implementation position; recheck before contract or deployment sign-off.

The AI Act entered into force on 1 August 2024. Prohibited-practice rules and AI literacy have applied since 2 February 2025, and governance and general-purpose AI obligations since 2 August 2025. Regulation (EU) 2026/1744 (Digital Omnibus on AI) entered into force on 27 July 2026. Most remaining AI Act provisions apply from 2 August 2026, while Chapter III high-risk rules apply from 2 December 2027 for Annex III systems and 2 August 2028 for Annex I systems. These dates are now set in the amended Regulation.

1 Aug 2024	2 Feb 2025	2 Aug 2025	2 Aug 2026	2 Dec 2027	2 Aug 2028
AI Act entered into force	Prohibited practices + AI literacy apply	Governance + GPAI obligations apply	Most remaining AI Act provisions apply	Annex III high-risk rules apply	Annex I product-safety route applies

Draft guidance status

The Commission's high-risk classification guidelines remain drafts as of 2 August 2026. A targeted consultation closed on 23 July 2026 and final adoption is pending. The drafts are useful but non-binding; authoritative interpretation remains with the CJEU and competent authorities.

Potential AI Act position for Aware Mate

- Aware Mate is an AI system within the scope of the AI Act.
- It is not designed for prohibited practices. In particular, it is not an emotion-recognition system for workplace use and does not perform identity recognition, social scoring or automated adverse employment action.
- Classification remains deployment-specific. The most relevant Annex III route is worker-management monitoring or evaluation if a deployment turns Aware Mate into person-level employer monitoring, scoring, profiling, HR evidence or performance evaluation.
- An aggregate-first safety-support configuration may fall outside the Annex III worker-management route only where its intended purpose and actual use do not involve monitoring or evaluating workers for employment-related purposes. Deployment-specific legal review remains required.
- ELNAV.AI maintains conservative documentation, logging, human-oversight and post-market controls aligned with potential high-risk obligations, without claiming a classification outcome before deployment-specific assessment.

3. 2026 classification guidance and practical impact

The AI Omnibus now fixes the high-risk timeline and clarifies the safety-component definition. The 2026 classification guidelines remain draft and continue to put weight on intended purpose, the full product presentation and whether outputs monitor or evaluate natural persons. Documentation and deployment controls therefore remain central.

Issue	Guidance signal	Aware Mate control
Intended purpose	Provider materials, instructions, sales language and technical documentation must be consistent. A simple exclusion in terms is not enough if the product is promoted for high-risk use.	Use one controlled intended-purpose statement across deck, one-pagers, SOW, DPA, product data sheet and website.
Human involvement	Human-in-the-loop design does not by itself remove high-risk classification; it is a control and, for high-risk systems, a compliance requirement.	Describe advisory design as a mitigation, not as an automatic low-risk classification.
Worker-management route	Monitoring or evaluating worker performance or behaviour can be an autonomous high-risk category. Safety support may fall outside this route only where intended purpose and actual use do not involve employment-related monitoring or evaluation.	Keep local support, aggregate-first reporting, no individual rankings, no HR use and no manager-facing live officer status.
Biometrics and emotion	The AI Act biometric-data concept is broader than GDPR. Emotion recognition means identifying or inferring emotions or intentions from biometric data.	No identity recognition, no sensitive categorisation, no emotion, stress, mood or intent inference. Fatigue/physical-state cues are treated as safety signals only.
Filter mechanism	Article 6(3) may exempt some Annex III systems, but the conditions are narrow and cannot apply where profiling is performed.	Do not rely on the filter in external claims unless a deployment-specific assessment is recorded and legal review supports it.
Annex I / safety component	The Omnibus clarifies that a safety component fulfils a safety function where its intended purpose is to prevent or mitigate health or safety risks. Article 6(1) also requires an Annex I product route and the relevant third-party conformity-assessment condition.	Do not treat 'advisory' as conclusive. Keep optional I/O non-interfering and reassess classification if a regulated product relies on Aware Mate or the intended purpose changes.

Classification trigger map

Default / lower classification pressure	Local bridge prompting; aggregate day or period trends; no person-level outputs; no raw video ashore by default; no HR, discipline or fitness use.
Legal review before activation	Exact event timestamps linked to watch schedules; event-level shore dashboard; optional raw-video buffer access; small-crew indirect identification; BNWAS-adjacent integration beyond non-interfering dry contact.
Avoid unless separately designed and reviewed	Officer scorecards, individual alert histories for management, live officer status, personnel ranking, emotional-state labels, automated HR/discipline/fitness decisions, navigation-control integration.

4. Roles, obligations and training

ELNAV.AI as provider • Maintain controlled intended-purpose, excluded-use and instruction-for-use wording. • Maintain technical documentation, change-control, issue logging and corrective-action paths appropriate to the system and updates. • Provide human-oversight information, boundary conditions, safe-fallback guidance and integration boundaries. • Support deployment-specific classification assessment without replacing customer legal advice. • Maintain cybersecurity and controlled-update processes, including signed binaries or configuration control where applicable.	Owner / operator as deployer and usually controller • Confirm intended purpose, deployment configuration and whether any high-risk route is triggered. • Determine GDPR lawful basis and, where required, Article 9 condition, DPIA and LIA approach. • Assign competent persons for human oversight, monitor operation, and act on identified risks or serious incidents. • Approve any remote support, additional export or event-level review in writing and reflect it in the DPA/SOW. • Inform affected employees and, where applicable, workers representatives before workplace deployment.	
Integrator / OEM caution Installation partners should stay within the approved interface and installation scope. Rebranding, substantial modification, safety-component integration or changes that alter intended purpose may trigger a fresh provider-obligation and classification review.		
AI literacy remains applicable Following the Omnibus amendment, providers and deployers must take measures to support role-appropriate AI literacy. Training should cover intended use, boundaries, failure modes, alert handling, override / escalation paths, data-handling rules and anomaly reporting.		
Workstream	Typical evidence / document	Current external-safe position
Intended purpose + instructions	System boundary, intended use, excluded uses, operator guidance	High-level wording can be shared externally
Classification assessment	Pilot/deployment assessment template and decision record	Template can be shared; customer-specific conclusion may remain controlled
Human oversight SOP	Pause, override, appeal, escalation and nuisance-alert flow	Template can be described externally; customer SOP stays controlled
Post-market monitoring	Issue log, change log, incident handling, corrective-action path	Policy can be summarised externally; records remain controlled
Data protection	Data handling summary, DPA, retention/access model, DPIA/LIA support	Summary can be external; signed DPA/DPIA remain controlled
Security controls	RBAC, audit logging, hardening, update policy	Only high-level summary should be external

5. GDPR and data-protection position

The GDPR analysis remains deployment-specific. The wording below is a conservative position for pilots and live deployments, not a substitute for customer legal advice.

Controller / processor split	The vessel owner or operator will usually be the controller for crew / bridge personnel data and onboard procedures. ELNAV.AI will usually act as processor for agreed support or reporting flows.
Personal-data / biometric question	Video-derived analysis of facial cues and behavioural indicators may involve personal data and may also engage biometric-data considerations, even where the system is not used to identify individuals. The AI Act concept of biometric data is broader than the GDPR concept.
Lawful basis	The controller must determine the Article 6 lawful basis and, where special-category rules are engaged, the Article 9 condition. A DPIA is strongly recommended before pilot or deployment in workplace settings.
Default data handling	By default, bridge workstation video is processed locally on the vessel edge unit. Standard off-vessel outputs are limited to aggregated or de-identified event metadata and review metrics; indirect-identification risk remains deployment-specific.
No profiling / scoring	The standard configuration should not create individual personnel profiles, officer rankings, person-level scores, management-facing individual histories or automated conclusions about performance or fault.
Indirect identification	Watch schedules can identify people indirectly. Avoid exact timestamps, time-of-day charts and small-cell reporting in basic aggregated review unless specifically justified and access-controlled.
Retention	Derived metrics and short event logs follow operator policy. Optional Critical buffers, if enabled, should be encrypted, access-controlled, logged and time-limited unless lawfully preserved.
Transfers / remote support	Any remote support, additional export or transfer outside the EEA should be specifically approved in writing and reflected in the DPA/SOW and transfer safeguards.
Default shore review may include	Escalated review only with written authority

- Period-level alert trends and severity mix.
- Valid monitoring time and pilot KPI summary.
- Alert load per monitored hour.
- Aggregated long-absence flags, if enabled.
- Context-weighted review, if agreed and available.

- Exact event timestamps where they can reveal watch assignments.
- Access to optional Critical buffers.
- Raw video export or remote diagnostic snippets.
- Incident review materials beyond aggregate reporting.
- Any review that may affect employment, discipline or fitness assessment.

6. Annex I, integrations and safety-component boundary

Aware Mate is currently positioned as an advisory support layer and is not intended to replace BNWAS or act as a mandatory navigational safety system. Under Regulation (EU) 2026/1744, the safety-component analysis must consider intended purpose and the consequences of failure; the advisory label alone is not decisive. No Annex I route has been established for the current standalone pilot configuration, but deployment-specific product-law review remains required before any regulated integration.

Question	Default answer	Action if YES / UNCLEAR
Is Aware Mate itself placed on the market as a regulated Annex I product?	Not established for the current standalone pilot configuration.	Escalate to Annex I product-safety and conformity-assessment review.
Is Aware Mate intended as a safety component of a regulated product that relies on it for safe operation?	Not established. The current configuration is standalone, advisory and non-navigation-control.	Apply the amended safety-component definition; review class/flag/sectoral rules and the relevant third-party conformity-assessment route.
Does Aware Mate replace BNWAS or a mandatory bridge safety system?	No. BNWAS and operator SMS remain independent.	Do not activate under standard scope. Separate certified configuration required.
Does optional I/O interfere with standard BNWAS or alarm-chain operation?	No. Optional dry contact should be non-interfering and fail-safe.	Require installer/OEM/customer approval, non-interference test and updated classification assessment.
Does Aware Mate control ECDIS, radar, autopilot, steering or COLREG decision-making?	No. This is outside intended purpose.	Fresh product, safety and legal review required before design or customer discussion.
Does a customer want event-level outputs for management or incident review?	Not by default.	Define purpose, lawful basis, access control, retention, worker notice and classification review before enabling.

Engineering controls for advisory integrations

- Document every integration as optional, customer-approved and non-interfering unless certified otherwise.
- Use fail-safe behaviour: loss of Aware Mate should not impair BNWAS or standard bridge operation.
- Test the alarm chain with Aware Mate connected and disconnected before acceptance.
- Record firmware, configuration, cabling, dry-contact logic and installer sign-off in the commissioning file.
- Treat any shift from advisory output to safety-component reliance as a material change requiring classification review.

7. Document classification and disclosure model

Can be shared in an external-safe version • System boundary, advisory-only role and privacy-by-design architecture. • Current AI Act timeline, AI Omnibus position and draft-guidance status. • High-level role allocation and need for deployment-specific legal confirmation. • High-level data handling, retention model and statement that DPIA/DPA work remains separate. • The high-level classification controls and trigger map in this overview.	Keep controlled / NDA • Customer-specific classification conclusions, legal memoranda, DPIAs, ROPAs and negotiated DPA terms. • Detailed technical file index, model card, acceptance thresholds and substantial-modification analysis. • Raw logs, selected clips, vessel-level diagnostic snippets and customer-specific incident materials. • Security hardening details, threat models, credentials, architecture maps and update procedures. • Commercial pricing, partner terms or vessel/crew-specific details unless approved for sharing.
--	---

Official and internal sources checked for this update

- Regulation (EU) 2024/1689 (AI Act) and the European Commission implementation timeline, accessed 2 August 2026.
- Regulation (EU) 2026/1744 (Digital Omnibus on AI), published 24 July 2026 and in force from 27 July 2026.
- European Commission, draft Guidelines for providers and deployers of AI high-risk systems; targeted consultation closed 23 July 2026 and final adoption remained pending on 2 August 2026.
- European Commission, AI literacy and implementation-support materials.
- European Commission, Guidelines on prohibited AI practices and Guidelines on the definition of an AI system, February 2025.
- Aware Mate Privacy Boundaries v1.2, Crew Trust Charter v1.1, Pilot SOW and commissioning templates.

Non-binding summary
This overview is for evaluation, pilot scoping, procurement and partner review. Customer-specific AI Act classification, GDPR lawful basis, transfer assessment and documentation requirements must be confirmed by the deployer/controller with appropriate legal guidance.