PUBLIC EXECUTIVE PREVIEW

AIGN CRITICAL OS

Governing Machine Authority in Critical Systems

A C-level preview of a proprietary governance and assurance architecture for AI operating in critical and high-consequence environments.

CRITICAL CONSEQUENCE	MACHINE AUTHORITY	RUNTIME CONTROL	DEFENSIBLE EVIDENCE

Based on the full AIGN Critical OS 1.4 reference work

This public preview presents selected executive principles only. The underlying 80+ page reference work contains the detailed classification criteria, control architecture, decision logic, assurance method, implementation artefacts and mapping structures and is not publicly distributed.

Patrick Upmann | Founder, AIGN.Global | Architect, AIGN Critical OS

Executive discussion paper. AIGN Critical OS is author-originated and is not a law, consensus standard, certification scheme or institutional endorsement.

WHY THIS MATTERS

Critical AI changes what an organisation must be able to govern

Enterprise AI governance commonly begins with inventory, legal classification, model risk, policy, documentation, human oversight and monitoring. Those disciplines remain necessary in critical environments - but they are not sufficient when AI can influence consequential action.

The additional problem is consequence

AI can influence service continuity, physical safety, public warning, operational control, financial stability, emergency response, cross-border coordination or the behaviour of other technical systems.

The shift is not simply from less capable AI to more capable AI. It is from systems that produce outputs to systems that can participate in organisational action.

From output to action

AI may communicate, prioritise, trigger, modify or control - not only recommend.

From access to authority

Technical permission can enable organisational effects without proving that those effects were legitimately delegated.

From static control to changing conditions

Authority that was appropriate at deployment may become inappropriate when dependencies, trust or operating conditions change.

For executives, the governance question therefore changes:

The management question

Can the organisation demonstrate what an AI system or agent was authorised to do, whose authority it exercised, whether that authority remained valid under changing operating conditions, and who remained accountable?

Capability is not authority

Capability is what a machine can do. Authority is what an organisation permits it to do.

In conventional software environments, access is often treated primarily as a technical permission issue. In agentic and critical AI, permissions can enable systems to communicate, trigger, modify or control on behalf of the organisation.

CAPABILITY	ACCESS	AUTHORITY	ACTION	CONSEQUENCE	ACCOUNTABILITY

AIGN design principle

A technically possible action should not automatically become an organisationally effective action. Authority must be deliberate, bounded and revocable.

The full reference work operationalises this distinction through proprietary classification, authority-boundary, trust, runtime-decision and evidence structures. Those structures are intentionally not reproduced in this public preview.

What AIGN Critical OS is - and what it is not

AIGN CRITICAL OS IS	AIGN CRITICAL OS IS NOT
An author-originated governance operating architecture for machine authority in critical systems.	A law, statutory rule, consensus standard or governmental certification scheme.
A vendor-neutral model for connecting consequence, delegated authority, changing trust conditions, runtime control and evidence.	A replacement for cybersecurity, safety engineering, privacy, operational resilience, GRC, IAM/PAM or AI assurance.
A management and assurance layer designed to interoperate with existing governance, risk and engineering regimes.	Automatic compliance, presumption of conformity or endorsement by NATO, NIST, ISO/IEC or another institution.
A practical basis for assessment and implementation in high-consequence environments.	A substitute for sector engineering, safety cases, legal interpretation or required certification.

Public boundary

This preview communicates the problem, the positioning, the applicability and selected executive examples. It does not disclose the detailed criteria, thresholds, decision rules, control requirements, assurance logic or implementation artefacts contained in the full reference work.

WHERE IT MATTERS

Critical AI is sector-neutral - consequence is not

AIGN uses 'critical systems' as an international umbrella term. Legal critical-infrastructure designations remain jurisdiction-specific.

Energy & utilities	AI supporting grid, process, restoration or optimisation decisions.
Digital infrastructure & telecom	AI and agents operating across networks, platforms and privileged workflows.
Banking, payments & financial infrastructure	AI influencing or executing consequential financial operations.
Health & medical	AI affecting clinical, operational or device-related decisions.
Transport & mobility	AI supporting routing, traffic, rail, aviation or autonomous operations.
Public warning & emergency services	AI contributing to detection, prioritisation, dispatch or crisis intelligence.
Manufacturing & chemicals	AI interacting with industrial, robotic or process-control environments.
Defence & dual use	AI supporting mission, logistics, sensing or autonomous functions.

The common governance question

Where does AI move from analysis into effective organisational action - and how can that authority remain defensible when consequence, trust or dependencies change?

ILLUSTRATIVE EXECUTIVE USE CASES

What this looks like in practice - without disclosing the control method

The following examples are deliberately high-level. They show where the governance problem appears; the detailed sector patterns, decision logic, controls and assurance criteria remain in the proprietary reference work.

SECTOR	ILLUSTRATIVE AI ROLE	POTENTIAL CONSEQUENCE	EXECUTIVE GOVERNANCE QUESTION
Energy & utilities	Forecasting, restoration or operational optimisation	Grid stability, service continuity, physical equipment	When may AI move from recommendation into bounded operational action?
Digital infrastructure & telecom	Autonomous network remediation and privileged operations	Wide-area outage and cross-sector cascade	What privileged actions remain legitimate when identity, network or provider trust degrades?
Banking, payments & financial infrastructure	Fraud response, payment routing, treasury or operational agents	Transaction, liquidity, market or systemic impact	Which actions may execute automatically - and within whose delegated authority?
Health & medical	Triage, clinical support, operational optimisation or device intelligence	Patient safety, delayed care and professional accountability	Which decisions can AI influence, and which must remain with competent human authority?
Transport & mobility	Routing, traffic management, fleet operations or autonomous functions	Safety, disruption and cross-network consequence	At what point does decision support become operational control?
Public warning & emergency services	Detection, prioritisation, dispatch and early warning	Population-scale false alarm, delay or misallocation	Who is the competent authority to issue, change or suppress a consequential warning?

Illustrative only. These examples are not sector certifications, legal conclusions or substitutes for sector engineering and local applicability assessment.

INTEROPERABILITY

Designed to sit across existing governance, risk and engineering regimes

AIGN Critical OS is designed as an interoperability layer, not a replacement regime. Existing disciplines and external requirements constrain the permissible authority space; AIGN focuses on the organisational question of what machine authority remains valid and defensible.

AI governance & risk	NIST AI RMF and related critical-infrastructure work; ISO/IEC 42001 family and AI risk / impact disciplines.
Cyber, OT & safety	IEC 62443, IEC 61508 and applicable sector safety families.
European regulatory landscape	EU AI Act, NIS2, CER, DORA, CRA and GDPR.
German critical infrastructure	German KRITIS cyber and resilience requirements.
Defence & allied context	NATO and allied public AI / autonomy principles.
Sector-specific regimes	Applicable product, safety, operational, privacy and professional requirements.

Reference boundary

AIGN does not claim equivalence, conformity, certification or institutional endorsement. 'Global' refers to a jurisdiction-neutral governance core designed for versioned sector and jurisdiction application - not a claim that one static document contains every law or standard in force everywhere.

The full reference work contains the underlying mapping and applicability structures. These are not reproduced in this public preview.

FROM REFERENCE WORK TO IMPLEMENTATION

AIGN Critical AI Authority Review

A focused executive assessment for organisations operating AI in critical or high-consequence environments.

Typical management outcomes

- A clear view of where AI already influences or exercises consequential organisational authority.
- Identification of priority authority, control, dependency and evidence exposure.
- A board-ready view of critical AI exposure without replacing existing cyber, safety, resilience or compliance programmes.
- A practical basis for deciding where deeper governance design, runtime control or implementation is justified.

The starting question

> If one of your AI systems acted incorrectly tomorrow, could you demonstrate what it was authorised to do, why that authority was still valid, which controls applied and who remained accountable?

Public preview vs. full reference work

PUBLIC PREVIEW	FULL REFERENCE WORK
Problem, executive principles, applicability, illustrative sectors and external reference boundary.	Detailed methodology, classifications, runtime logic, assurance criteria, implementation artefacts and mapping structures.

The full AIGN Critical OS reference work is not publicly distributed. It is used selectively in assessment and implementation engagements.

Patrick Upmann | Founder, AIGN.Global | Architect, AIGN Critical OS

For executive discussion, assessment and implementation enquiries.