AI-AWARE RESILIENCE GOVERNANCE™

Governing the Resilience of Mission Critical Business Services in an AI-Enabled Enterprise

An ORG Initiative Whitepaper

September 2026

Central Proposition

AI does not create a new need for organizational resilience governance. It changes the nature of the dependencies that resilience governance must understand.

Published by: Operational Resilience Governance Initiative™ September 2026

Executive Summary

Artificial intelligence is moving from a technology used to generate or analyze information toward AI-enabled capabilities that can recommend, decide, and act within organizational workflows. At the same time, AI is increasingly embedded inside enterprise products and third-party services. This creates a resilience challenge that is not adequately described by asking whether an organization "uses AI." The more consequential question is whether an AI capability has become a dependency of a Business Function supporting a Mission Critical Business Service—and whether changes in that AI dependency can alter the condition, behavior, decision authority, or action authority of the Business Function.

Operational Resilience Governance™ (ORG) provides a useful perspective because its purpose is not to govern a particular technology. ORG provides an executive governance context through which existing resilience disciplines, capabilities, information, evidence, and operational results can be understood, evaluated, and converted into accountable executive conclusions. The current ORG publications explicitly position ORG as a discipline that does not replace continuity, cybersecurity, risk, supplier, technology, or other disciplines.

This whitepaper therefore proposes an AI-aware interpretation of ORG rather than a separate AI governance framework. The objective is to determine whether the existing ORG concepts—Mission Critical Business Service, Business Function, Enterprise Support Service, Service Continuity Requirement, Governance Evidence, Required Resilience Condition™, Resilience Condition™, Executive Certification™, and Executive Certification Cascade™—are sufficient to govern resilience in an environment where dependencies may contain AI systems capable of changing behavior or taking actions.

The ORG Question
If an AI-enabled capability is a dependency of a Business Function supporting a Mission Critical Business Service, what condition must that capability be in, what evidence demonstrates that condition, and what does the accountable executive need to know before establishing Executive Certification™?

1. Purpose and Scope

This paper is a research and conceptual working paper. It is intended to support the continued development of the ORG discipline before revisions are made to the five sections of the two current ORG publications.

The paper does not establish a new ORG Canonical Object for AI, does not prescribe an AI control framework, and does not attempt to replace established AI governance, cybersecurity, model risk, data governance, supplier governance, or regulatory requirements. Those disciplines remain responsible for the work within their respective domains.

Instead, the paper examines where AI creates resilience implications that must become visible within the existing ORG governance context.

The analysis is organized around six questions:

- What is an AI System from the perspective of the AI community?
- How can AI become a dependency of a Mission Critical Business Service?
- Why are "decide" and "act" materially different from "inform" and "recommend"?
- How can AI dependencies change without an equivalent change in the customer's own systems?
- What evidence should an organization expect from vendors that embed AI?
- How can ORG remain technology-independent while becoming AI-aware?

2. The Existing ORG Foundation

The current ORG Discipline establishes ORG as an Executive Governance Discipline™ centered on Executive Certification™ of the resilience condition of Mission Critical Business Services. It explicitly states that ORG does not replace the disciplines that produce resilience capabilities, information, evidence, assessments, or operational results.

The current ORG Implementation Guide further describes implementation as a governance information transformation, not a software installation. Existing organizational information is mapped into the ORG Canonical Model, related to accountability, requirements, condition, evidence, and certification, and presented to accountable executives.

The Canonical Model establishes the Business Function as the principal bridge between a Mission Critical Business Service and the supporting capabilities required to sustain it. Enterprise Support Services may represent an application, infrastructure capability, platform, shared service, or other enterprise capability upon which Business Functions depend.

The current Canonical Model contains ten active canonical objects. Required Resilience Condition™, Certifiable Governance Object™, Executive Evaluation, and Executive Certification Cascade™ are represented as governance concepts or relationships rather than separate active objects

Existing ORG construct	Relevance to AI-aware resilience
Mission Critical Business Service Profile	Defines the business service whose resilience matters; AI is considered in the context of sustaining that service.
Business Function Profile	Provides the operational bridge where AI-enabled capabilities may support execution.
Enterprise Support Service Profile	Can represent an application, platform, shared service, or other capability that contains or depends upon AI.
Service Continuity Requirement	Defines continuity outcomes applicable to a Business Function; AI may affect the capability required to meet them.
Governance Object	Provides the governed object context without requiring a separate “AI object.”
Required Resilience Condition™	Can express the condition required of a governed capability, including relevant AI behavioral or authority characteristics.
Governance Evidence	Can incorporate AI evaluations, monitoring, vendor assurance, testing, change records, incidents, and other evidence.
Resilience Condition™	Represents current condition after evaluating available evidence against the required condition.
Executive Certification™	Provides the accountable executive's conclusion regarding the governed condition.
Executive Certification Cascade™	Connects certifications across dependency and accountability boundaries.

3. What Is an AI System?

NIST Trustworthy and Responsible AI – NIST AI 100-2e2025 defines an artificial intelligence system broadly. Its AI Risk Management Framework is intended for organizations designing, developing, deploying, or using AI systems, and emphasizes trustworthiness across the design, development, use, and evaluation of AI products, services, and systems.

For ORG purposes, the important implication is that an AI System should not be treated as synonymous with an AI model. An operational AI system may include an application, model, data, knowledge sources, orchestration, interfaces, tools, infrastructure, monitoring, and external providers. The exact architecture varies by product.

Component or characteristic	ORG relevance
Application / workflow	Where AI enters the Business Function and where business consequences occur.
AI model	Produces or transforms outputs; model changes can affect behavior without a customer application change.
Data / knowledge	Inputs and retrieved information can influence outputs and decisions.
Orchestration	Controls how models, data, tools, and workflow components interact.
Tools / APIs	Can move AI from information production toward operational action.
Infrastructure / provider	Availability and dependency concentration can affect the AI-enabled capability.
Monitoring / evaluation	Produces evidence about current performance, behavior, security, and changes.

The key ORG question is therefore not “How many AI systems does the organization have?” but **“Which AI-enabled capabilities have become material dependencies of Business Functions supporting Mission Critical Business Services?”**

4. Embedded and Potentially Hidden AI Dependencies

AI can be introduced by an organization directly, by an enterprise application provider, by a platform provider, or through a third-party component embedded within another service. NIST's Generative AI Profile explicitly notes that GAI value chains involve many third-party components, including datasets, pretrained models, and software libraries, and that extensive integration into other devices and services can make it difficult to attribute system behavior to any single source.

This creates a resilience visibility problem. A customer may map a Business Function to a vendor service while the vendor's internal architecture has evolved to include AI capabilities, external models, agents, tools, or other dependencies. The resilience map can remain unchanged even though the behavior or dependency characteristics of the service have changed.

Potential Visibility Gap

The customer may know: Business Function → Vendor Service. The vendor may know: Vendor Service → AI Capability → Model → Data/Knowledge → Tools/APIs → AI Provider → Infrastructure. AI-aware ORG must determine whether the additional dependency information is material to the resilience condition of the Business Function.

5. Inform, Recommend, Decide, Act

One of the most important AI-related resilience distinctions is the degree of operational authority assigned to an AI-enabled capability.

AI role	Typical pattern	Potential resilience significance
Inform	AI produces information for a person.	Human remains responsible for interpretation and action; AI error can still affect outcomes.
Recommend	AI proposes an action or decision.	Human review remains a control point, but recommendation quality becomes important.
Decide	AI determines an outcome within defined authority.	Incorrect reasoning or decision can directly affect the Business Function.
Act	AI or an AI-enabled system executes the outcome.	Incorrect decisions can propagate directly into operational execution.

The transition from “recommend” to “decide” to “act” is not simply a technology enhancement. It can change the resilience characteristics of the dependency. The organization may need assurance not only that the capability is available and functioning, but that it remains within its authorized decision and action boundaries.

This issue is especially relevant to AI agents. NIST's current adversarial machine learning work recognizes attack and misuse concerns for AI systems and notes that the consequences of attacks become more significant in high-stakes domains.

6. The Resilience Problem of a Wrong Decision

Traditional resilience discussions often emphasize loss of availability: a system is unavailable, a facility is lost, a supplier fails, or infrastructure becomes inaccessible. AI introduces another class of concern: the system may be available and technically functioning while producing an unacceptable decision or action.

Condition	Traditional interpretation	AI-aware interpretation
Available	System is accessible.	AI capability is accessible and responsive.
Performing	System executes its function.	AI produces outputs or decisions.

Correct	Rules/processes produce expected results.	Outputs and decisions remain within required behavioral parameters.
Authorized	System performs defined functions.	AI remains within its delegated decision/action authority.
Secure	System is protected against unauthorized access.	AI inputs, models, knowledge, tools, and interactions are protected against manipulation.

NIST AI 100-2e2025 identifies adversarial machine learning as a distinct security concern and describes attacks including evasion, poisoning, privacy, and misuse attacks across AI/ML lifecycles.

From an ORG perspective, the critical question is the business consequence: if an AI system is manipulated, produces an incorrect result, or takes an incorrect action, does the affected Business Function remain within its Required Resilience Condition™?

7. Vendor Responsibility and AI Change

A vendor that embeds AI into a customer-facing service creates an information asymmetry: the vendor may know the internal AI architecture while the customer sees only the service. The customer therefore needs sufficient information to determine whether a material AI change can affect a Business Function on which a Mission Critical Business Service depends.

NIST's Generative AI Profile recognizes the risks associated with third-party components and integration and encourages organizations to manage risks across the AI value chain.

This does not mean the vendor must disclose proprietary model details. The resilience need is more targeted: the customer needs to know when a change can materially affect capability, behavior, decision authority, action authority, dependencies, or the evidence supporting continued reliance.

Proposed Customer Resilience Question

Has AI been incorporated into, materially changed within, or removed from any capability, function, or service you provide to us that could affect the performance, continuity, decision-making, or execution of a Business Function supporting a Mission Critical Business Service?

If the answer is yes, the customer should seek information about:

- Where AI is used in the service.
- What the AI does: inform, recommend, decide, or act.
- What decisions or actions it is authorized to perform.
- What human oversight or intervention is required.
- What material models, providers, data, knowledge sources, tools, and external dependencies are involved.
- How material AI changes are identified, evaluated, tested, approved, and communicated.
- What evidence demonstrates that the AI-enabled capability remains within its intended and authorized operating condition.

- What contingency or fallback arrangements exist if the AI capability becomes unavailable or unacceptable.

8. AI Security Is Part of Resilience Evidence

AI security cannot be separated from resilience when an AI system is a material dependency of a Business Function. NIST AI 100-2e2025 explains that the statistical and data-based nature of machine learning creates attack vectors beyond traditional software, including manipulation of training data, adversarial inputs, and malicious interaction with models and their accessible information.

ORG does not need to become the cybersecurity discipline for AI. Cybersecurity, AI security, model security, privacy, and other specialists should produce the relevant evidence. ORG needs to determine whether that evidence is sufficient to establish the current condition of the governed Business Function or capability.

Existing discipline	Potential evidence available to ORG
Cybersecurity / AI security	Security assessments, vulnerabilities, incidents, attack testing, control status.
AI governance	Use-case approval, policy compliance, risk assessments, governance decisions.
Model validation / evaluation	Performance evaluations, robustness testing, model version assessments.
Data governance	Data quality, provenance, access, integrity, freshness.
Vendor management	Contractual commitments, service levels, change notices, third-party assurance.
Continuity / resilience	Recovery arrangements, alternate procedures, testing results, impact analysis.
Operations	Monitoring, incidents, performance, operational exceptions, change history.

9. What "AI-Aware ORG" Means

AI-aware ORG does not mean creating a parallel AI governance framework. It means ensuring that the existing ORG discipline can recognize and govern resilience implications when AI becomes part of a dependency.

Current ORG question	AI-aware extension
What Business Functions support the MCBS?	Which of those functions depend on AI-enabled capabilities?
What supporting capabilities exist?	Which contain or depend on AI systems?
What condition is required?	Does the required condition include behavioral, decision, or action boundaries?
What evidence exists?	Does the evidence demonstrate AI availability, integrity, performance, behavior, authority, and change condition where material?

What is the current condition?	Is the AI-enabled capability within the required condition, including material uncertainty?
What must the executive certify?	Can the executive reasonably rely on the evidence supporting the AI-dependent capability?
What changes matter?	Can a vendor/model/AI change alter the Business Function's resilience condition without a customer-initiated system change?

10. An AI-Aware ORG Governance Flow

A practical conceptual flow is:

Step	AI-aware ORG activity
1. Identify	Identify AI-enabled capabilities supporting Business Functions associated with Mission Critical Business Services.
2. Map	Map the AI-enabled capability to the Business Function and applicable Enterprise Support Service relationships.
3. Characterize	Determine whether AI informs, recommends, decides, or acts; identify material dependencies and authority boundaries.
4. Define condition	Determine what condition is required for the AI-enabled capability to support the Business Function.
5. Gather evidence	Use existing AI, cybersecurity, vendor, data, testing, monitoring, continuity, and operational evidence.
6. Assess	Evaluate evidence against the Required Resilience Condition™.
7. Establish condition	Establish the current Resilience Condition™ and material uncertainty.
8. Certify	Support accountable Executive Certification™ where required.
9. Cascade	Communicate the certified conclusion through the Executive Certification Cascade™.
10. Maintain	Re-evaluate when material AI, vendor, model, data, authority, or operational conditions change.

This approach is consistent with the current ORG implementation principle: existing organizational information is mapped to canonical meaning, related to governance, evaluated against condition, and used to support Executive Certification™.

11. Required Resilience Condition™ in an AI-Enabled Environment

The introduction of AI does not automatically require a new type of resilience condition. It may, however, require a more complete statement of the condition required for a Business Function.

For an AI-dependent capability, the required condition may include, where material:

- Availability and recoverability.
- Performance and capacity.
- Integrity of inputs and relevant knowledge.
- Accuracy or reliability appropriate to the use case.
- Behavioral stability within defined operating parameters.
- Defined decision authority.
- Defined action authority.
- Human oversight and intervention requirements.
- Security and resistance to material manipulation.
- Known and governed dependencies.
- Change notification and change evaluation.
- Fallback or alternate operating capability.

These are not proposed ORG Canonical Model fields at this stage. They are candidate characteristics for further research to determine whether they should be represented as requirements, evidence attributes, assessment criteria, or other governance information.

12. Executive Certification in the AI Era

The current Executive Governance Standard states that executive management should be able to understand what matters to the Mission, what condition is required, what the organization knows, what available evidence demonstrates, what condition the organization is currently in, what remains uncertain, what is being done about material conditions, and what requires executive action.

AI strengthens rather than weakens this requirement. An executive should not be asked to certify that an AI model is “safe” or “trustworthy” in the abstract. The executive should be able to evaluate whether the Business Function remains within its required resilience condition given the AI capabilities and dependencies on which it relies.

AI-Aware Executive Certification
The accountable executive should have sufficient Governance Evidence to determine whether an AI-dependent Business Function remains within its Required Resilience Condition™, including any material limitations, uncertainty, delegated decision authority, action authority, or dependency changes.

13. Executive Certification Cascade™ and AI Dependencies

The current ORG Standard explains that a Mission Critical Business Service depends on Business Functions, which depend on supporting capabilities, which may in turn depend on additional organizational capabilities,

Enterprise Support Services, and external or third-party services. The Executive Certification Cascade™ connects governance accountability, evidence, and Operational Resilience Posture™ through Executive Certification™.

AI can introduce another dependency layer without changing that governance relationship. An AI-enabled capability may be supported by an application owner, AI owner, model provider, data owner, supplier, or other accountable parties. ORG's question is not whether every technical participant becomes an executive certifier. It is whether the accountability and evidence needed to support the Business Function's certification remain traceable.

Possible dependency level	Potential governance evidence
AI-enabled capability owner	Capability performance, authorized behavior, operational monitoring.
Application / platform owner	Application condition, integration, fallback, change management.
AI provider / vendor	Model/service changes, material dependencies, testing/evaluation, incidents, service commitments.
Business Function executive	Conclusion regarding the function's required condition.
MCBS executive	Certification supported by lower-level evidence and certifications.
Enterprise executive	Enterprise-level conclusion about material resilience conditions and actions.

14. Why ORG Should Not Become Another AI Governance Framework

There are already significant AI governance and risk resources. NIST's AI RMF is intended as a voluntary, cross-sector resource for managing AI risks; its Generative AI Profile addresses risks associated with third-party components and integration; and NIST's 2025 adversarial machine learning publication establishes a security-focused taxonomy and terminology.

ISO/IEC 42001 provides a management-system foundation for AI governance, and ISO states that certification to ISO/IEC 42001 is voluntary and performed by independent certification bodies. cite turn0search5

These resources can produce valuable Governance Evidence for ORG. ORG's distinctive contribution is different: connect the evidence to the resilience of Mission Critical Business Services and to accountable executive certification.

Discipline / framework	Primary concern	ORG relationship
AI governance	Responsible and controlled use of AI.	Potential source of governance evidence.
AI security / adversarial ML	Protection against manipulation and attack.	Potential source of security and resilience evidence.

Model evaluation / validation	Model performance and behavior.	Potential evidence for Required Resilience Condition™.
Cybersecurity	Security of systems, data, interfaces, and infrastructure.	Potential evidence.
Vendor / supplier governance	Third-party risk and contractual assurance.	Potential evidence and dependency information.
Business continuity / resilience	Continuity and recovery of capabilities.	Existing resilience capability producing evidence.
ORG	Executive understanding and certification of MCBS resilience condition.	Integrates relevant evidence without replacing the source disciplines.

15. A Proposed Vendor Resilience Inquiry for AI-Enabled Services

The following is a proposed research artifact, not yet an ORG standard requirement. It illustrates the type of inquiry that could help close the visibility gap.

- Identify any AI-enabled capability incorporated into the product or service provided to the customer that could affect a Business Function supporting a Mission Critical Business Service.
- For each capability, identify whether AI informs, recommends, decides, or acts.
- Identify material decisions or actions that the AI system is authorized to perform.
- Identify material AI models, external AI providers, data sources, knowledge sources, tools, APIs, and other dependencies relevant to the capability.
- Describe how material changes to AI behavior, model, decision authority, action authority, or material dependencies are identified and communicated.
- Describe the evidence available to demonstrate continued operation within intended and authorized parameters.
- Describe material failure, fallback, human intervention, and contingency arrangements.
- Identify any material limitations or uncertainty that the customer should consider when relying on the capability.

16. Implications for the ORG Publications

This paper does not recommend immediate rewriting of the five ORG sections. Instead, it identifies areas that should be tested against the existing discipline before changes are made.

ORG section	Question for future review
Resilience Universe Map™	Does the Resilience Universe adequately recognize AI-enabled capabilities, AI supply chains, and dynamic dependencies within the existing environment?
Executive Governance Standard™	Do the governance requirements adequately address changes in dependency behavior, decision

	authority, action authority, and AI-related uncertainty?
ORG Implementation Guide™	Does implementation guidance explain how AI-related information and evidence can be mapped without turning ORG into an AI inventory or technical repository?
ORG Canonical Model™	Are existing objects and relationships sufficient to represent AI-related resilience information, or is additional canonical structure actually necessary?
ORG Canonical Specifications™	Can an ORG-enabled application represent and assess AI-related evidence, conditions, changes, and dependencies without prescribing an AI technology architecture?

The current Canonical Model already provides a technology-independent architecture centered on Mission Critical Business Services, Business Functions, Enterprise Support Services, Service Continuity Requirements, Governance Objects, Resilience Condition, Executive Certification, accountability, and lifecycle.

17. Recommended Research Agenda Before Publication Revision

- Establish a common AI terminology set using authoritative sources such as NIST and other major standards bodies.
- Map a real enterprise AI ecosystem, using Epic or another major enterprise platform as a case study, while distinguishing public documentation from inferred architecture.
- Identify the resilience implications of AI systems that inform, recommend, decide, and act.
- Study AI-specific dependency characteristics: model changes, behavior changes, provider concentration, data/knowledge dependencies, tools, APIs, and agent autonomy.
- Study vendor disclosure and change-management practices for embedded AI.
- Map AI security and adversarial ML risks to Business Function consequences rather than treating them solely as technology risks.
- Test whether the current ORG Canonical Model can represent these concerns without creating unnecessary new objects.
- Use the ORG POC to test representative AI-aware governance scenarios before changing the Canonical Specifications.
- Only after these tests, determine whether changes to the five publication sections are warranted.

18. Conclusion

AI is not necessarily a reason to create another resilience framework. It is a reason to examine whether the existing resilience governance discipline can see and govern a changing dependency environment.

The central issue is not whether an organization uses AI. It is whether AI has become part of the capability required to sustain a Mission Critical Business Service—and whether the behavior, authority, dependencies, or condition of that AI capability can affect the Business Function.

The most consequential transition may be from AI that informs or recommends to AI that decides and acts. Once AI participates directly in execution, resilience must consider not only availability and recovery but also whether the capability continues to operate within the behavioral and authority boundaries required by the Business Function.

ORG is well positioned to address this problem without becoming an AI governance framework. Its existing discipline is designed to connect Mission Critical Business Services, Business Functions, supporting capabilities, requirements, evidence, current condition, accountability, and Executive Certification™. The current implementation approach deliberately uses information produced by existing disciplines rather than replacing them.

Proposed ORG Perspective

ORG does not govern AI as a technology. ORG governs the resilience of Mission Critical Business Services in an environment increasingly enabled by AI. When AI becomes a material dependency of a Business Function, its relevant condition, behavior, authority, dependencies, changes, and evidence become part of the executive resilience question.

The next stage of the ORG Initiative should therefore be disciplined investigation rather than immediate expansion of the ORG model. The objective is to determine whether AI reveals a gap in the existing governance representation—or whether the existing ORG discipline is already sufficiently general, provided that it is applied with an AI-aware understanding of dependencies and condition.

References and Research Sources

- **NIST.** Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, 2023. https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10
- **NIST.** Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1, 2024; updated resource information 2026. https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence
- **NIST.** AI 100-2e2025, Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations, March 2025. https://csrc.nist.gov/pubs/ai/100/2/e2025/final
- **NIST.** AI Research — Security and Resilience; AI RMF Trustworthy AI in Critical Infrastructure Profile development. https://www.nist.gov/artificial-intelligence/ai-research-security-and-resilience
- **ISO.** ISO/IEC 42001 explained; certification is voluntary and performed by independent certification bodies. https://www.iso.org/home/insights-news/resources/iso-42001-explained-what-it-is.html

Working paper status: Research and conceptual development — not an ORG standard.

Published by: Operational Resilience Governance Initiative™ September 2026