RESILIENCE & RECOVERY

Staying operational, staying sovereign

The 15 Guardrails of AI — Pillar 5 of 5

A Vertex11 Whitepaper

2 guardrails · 8 controls · 4 key

G13 Operational Resilience, Continuity & Vendor Risk (4 controls, 2 key)

G14 Sovereignty & Continuity of Access (4 controls, 2 key)

Validated against NIST AI RMF 1.0, the EU AI Act (as amended by the 2026 Digital Omnibus), ETSI TS 104 223, and ISO/IEC 42001:2023

Executive summary

Every previous pillar reduces the incidents Resilience & Recovery has to handle. This pillar handles the ones that get through anyway — the model that fails, the vendor that withdraws, the jurisdiction that acts. It is the pillar that provides continuity: the ability to keep running, and to keep control of what you built, when the ground under the AI moves without warning.

The pillar contains two guardrails: Operational Resilience, Continuity & Vendor Risk, which ensures AI capabilities survive provider failure, degradation and lock-in; and Sovereignty & Continuity of Access, which ensures the organisation retains control over where its AI runs, whose law governs it, and whether it can move it or take it back. Together they carry 8 control statements, 4 of them key, validated against four international standards: NIST AI RMF 1.0, the EU AI Act (as amended by the 2026 Digital Omnibus), ETSI TS 104 223, and ISO/IEC 42001:2023. Every control traces to at least one standard provision.

This paper details the risks each guardrail counters, what the controls cover, and what is required across people, process, and technology to sustain them.

Why resilience and recovery is a distinct pillar

AI capability sits on a stack of dependencies most enterprises have never risk-assessed as continuity risks. Foundation-model providers, hosting jurisdictions, third-party feeds, model licences and cloud-tenanted inference — these are the load-bearing components of many AI-enabled services, and each carries a failure mode traditional IT disaster recovery was not designed to catch. Resilience & Recovery treats the model as a supplier relationship with jurisdictional exposure, not a piece of software with a backup.

Two questions define the pillar. *What happens when the AI capability we depend on becomes unavailable, degraded, or withdrawn?* And *do we still control where our AI runs, whose law governs it, and whether we can move it or take it back?* The first is a continuity discipline; the second is a sovereignty discipline. Both are load-bearing, and both are typically under-invested until exercised in anger.

The standards validation reinforced the positioning. NIST's MANAGE function requires recovery and continuity capability calibrated to risk. The EU AI Act's provider and deployer obligations extend into the supply chain, making vendor risk a governance question rather than a procurement one. ETSI TS 104 223 anchors vendor lifecycle risk explicitly. ISO/IEC 42001 embeds continuity and supplier relationships in the management system. Four standards, one message: an AI capability you cannot sustain through provider, technology or jurisdictional shock is not a capability — it is a bet.

The risks that manifest without it

Provider failure or withdrawal

The foundation model your product depends on is deprecated at short notice, or the provider changes terms in a way that makes continued use impossible. Capability disappears not because it failed, but because the party supplying it decided to stop supplying it — and there is no fallback because none was designed in.

Vendor lock-in and cost drift

Capability becomes hostage to a single provider's pricing, roadmap and licensing. What was a competitive market position at signature turns into a captive one over time, with no realistic exit path and costs that drift upward against a floor the organisation has no way to negotiate against.

Silent degradation

The model is not withdrawn — it is quietly changed under the same name. Performance drops, behaviours shift, and the system that appeared to work continues to appear to work while its outputs degrade. Degradation-in-place is often worse than outright failure, because it does not trigger the reflexes that a hard outage does.

Ecosystem concentration

A handful of foundation-model providers underpin much of the market. Concentration risk that would be flagged in any other critical supplier category is normalised in AI because the alternatives are few and the switching cost is real. Systemic exposure is treated as inevitable rather than assessed.

Sovereignty exposure

Data, compute or model residency sit in a jurisdiction that can, at short notice, alter the terms of the organisation's access — through export control, sanctions, legal process or policy change. The default choice of provider often carries a sovereignty posture the organisation never consciously accepted.

External compulsion or access removal

A third party — regulator, government, provider or court — removes the organisation's access to its own capability or data. Without arrangements to retain access despite such action, the organisation discovers that critical capability was contingent all along.

No exit path

Capability cannot be repatriated or migrated without unacceptable loss — of data, of tuning, of integration, of institutional knowledge. Exit provisions that were not negotiated at entry are rarely available at exit; portability tooling that was not validated in advance is rarely usable under pressure.

Untested continuity

Continuity plans exist on paper but have never been exercised against realistic provider-outage or jurisdictional-action scenarios. Their gaps are discovered at the moment of failure — which is the moment they matter most and can least be repaired.

Guardrail 13: Operational Resilience, Continuity & Vendor Risk

4 controls · 2 key | Pillar: Resilience & Recovery | Applicability: Universal

G13 governs *the ability to sustain critical AI capability through provider, technology or supply-chain shock*. It is the continuity half of the pillar.

What the controls cover

Critical AI dependencies — models, providers, platforms, feeds — are identified and their concentration risk assessed. The exercise reaches deeper than the direct supplier: a model accessed via a reseller is still exposed to the underlying provider's failure, and a hosting choice that shares infrastructure with a competing service inherits that service's incident profile.

Continuity arrangements exist for critical AI capability, aligned to defined recovery objectives. What constitutes an arrangement depends on how the capability is used and how quickly it must return: fallback to an alternative provider, safe-mode degradation to a reduced-capability path, or alternative provision from a different technology stack. The choice is deliberate and matched to criticality, not defaulted to whichever the incumbent offered as an upsell.

Vendor risk is assessed and managed across the provider lifecycle — viability, lock-in, exit friction, service change, licence change. This is a lifecycle discipline, not a procurement one: the vendor risk that matters is the one that emerges after signature, when the switching cost has been paid and the balance of power has moved. And resilience arrangements are tested, not merely documented. An untested continuity plan is a hope, not a control; discovery of its gaps at the moment of failure is the classic mode.

Guardrail 14: Sovereignty & Continuity of Access

4 controls · 2 key | Pillar: Resilience & Recovery | Applicability: Universal

G14 governs *whether the organisation retains control over where its AI runs, whose law governs it, and whether it can move it or take it back*. It is the sovereignty half of the pillar — a discipline distinct from continuity, and increasingly separable from it.

What the controls cover

Where AI runs, where its data resides, and under whose authority it operates are known and deliberately controlled — not accepted as vendor defaults. Provider region, sub-processor list, key management arrangements and applicable jurisdictional law are properties the organisation records, reviews and can defend.

Arrangements exist to retain access to critical capability and data despite external compulsion, withdrawal or jurisdictional action. The specific mechanism varies by workload and risk posture: sovereign or contractually-protected placement, encryption-at-rest with customer-held keys, mirrored replicas across jurisdictions, or contractual guarantees enforceable in the relevant courts. What does not vary is that the arrangement exists and has been validated.

Exit and portability provisions allow capability to be migrated or repatriated without unacceptable loss — of data, of tuning, of integration or of institutional knowledge. Exit is negotiated at entry, because it is

not available at exit; portability tooling is validated in advance, because it is rarely usable under pressure. Sovereignty-critical workloads are identified and placed consistent with the organisation's risk and jurisdictional posture, rather than distributed by default across whichever regions the provider found convenient.

Building the capability: people, process, technology

People

Role	Responsibility
AI governance lead	Accountable for continuity and sovereignty posture across the AI estate; owns the recovery-and-improvement loop that feeds R&R findings back into use-case approval.
Business-unit / system owners	Declare criticality, recovery objectives and acceptable degradation modes per system; approve the continuity arrangement that applies.
Enterprise architecture	Maps dependencies across models, providers, platforms and feeds; designs fallback, safe-mode and alternative-provision paths proportionate to criticality.
Procurement & vendor management	Owns vendor lifecycle risk — viability, lock-in, service change, licence change, exit friction — and secures portability and exit provisions in contract before signature.
Legal & regulatory affairs	Jurisdictional analysis; sanctions and export-control posture; enforceability of contractual guarantees in the relevant courts.
Security & platform engineering	Implements sovereign hosting patterns, customer-held keys, cross-jurisdiction mirroring and access-continuity primitives.
Business continuity / DR	Integrates AI recovery into wider BC/DR; runs continuity exercises including provider-outage and jurisdictional-action scenarios.

The defining accountability question in this pillar is who owns the vendor and jurisdictional posture as a whole — not the individual contracts, but the concentration risk, the exit readiness and the sovereignty stance across the estate. Left as a distributed procurement question, that posture drifts toward whichever provider is easiest at each moment of choice, and the aggregate exposure only becomes visible under stress.

Process

Dependency mapping. Every critical AI capability is traced to its underlying models, providers, platforms, feeds and hosting jurisdictions — including indirect dependencies via resellers and downstream infrastructure.

Concentration risk assessment. Provider, model and jurisdiction concentration across the AI estate is assessed as a whole, not per contract, and reviewed on cadence and on material market change.

Continuity design proportionate to criticality. Recovery objectives set per system determine whether fallback, safe-mode or alternative-provision applies — a deliberate design choice, not a defaulted one.

Sovereignty posture per workload. Placement, key management and applicable-law choices are made against the organisation's declared risk and jurisdictional posture, not against provider convenience.

Exit and portability in contract before signature. Data portability, model portability where possible, and clean-exit provisions are negotiated at entry, not sought at exit when the balance of power has moved.

Continuity exercises on cadence. Realistic provider-outage and jurisdictional-action scenarios exercised end-to-end, producing evidence that the arrangement works rather than that it exists.

Vendor lifecycle review. Viability, service change, licence change and exit friction reviewed across the provider lifecycle — not just at procurement.

Technology

Dependency inventory — records provider, model, jurisdiction, licence terms, sub-processors and criticality for every AI capability, and is refreshed on change.

Fallback and safe-mode primitives — at model and orchestration layers, with routing that can switch providers, degrade gracefully or invoke alternative-provision paths on defined triggers.

Multi-provider abstraction — where continuity requires it, with interfaces designed so the substitution is operational, not aspirational.

Sovereign hosting patterns — customer-held keys, region-pinned inference, mirrored replicas across jurisdictions, with the choice matched to workload sensitivity.

Portability tooling — model weights (where licensed), prompts, evaluations, tuning artefacts and integration configurations exportable in a form the receiving system can actually use.

How the standards shaped this pillar

NIST AI RMF 1.0 requires recovery and continuity capability calibrated to mapped risk, and treats supply-chain risk as first-order rather than a procurement footnote.

The EU AI Act (as amended by the 2026 Digital Omnibus) extends provider and deployer obligations into the supply chain — meaning vendor selection, allocation of AI-related responsibilities and continuity of compliance are governance questions, not just commercial ones.

ETSI TS 104 223 anchors vendor risk lifecycle management explicitly, including dependency management and the operational-security implications of third-party provision.

ISO/IEC 42001 embeds continuity planning and supplier-relationship management within the management system, connecting R&R findings back to the policy, roles and management review that Assurance operates.

The net effect: the Resilience & Recovery pillar covers continuity of capability, control of jurisdiction and readiness to exit, to the expectations of all four standards, from a single set of controls.

A maturity path

Level 1 — Ad hoc. Dependencies partially known. No continuity design for AI capability specifically. Sovereignty accepted as vendor default. No exit provisions in contract. Continuity untested.

Level 2 — Documented. Dependency inventory exists. Concentration risk assessed at least annually. Continuity arrangements defined for critical systems. Exit and portability provisions in new contracts. Sovereignty posture stated per workload.

Level 3 — Enforced and tested. Continuity exercised on cadence, including provider-outage scenarios. Sovereign placement implemented for identified workloads. Portability tooling validated. Vendor lifecycle risk reviewed beyond procurement.

Level 4 — Continuously assured. Multi-provider abstraction where warranted. Sovereignty posture verified against jurisdictional-risk changes in-quarter. Exit rehearsed for critical vendors. Concentration risk actively managed at the estate level, not just recorded.

The goal is not uniform maturity across both guardrails on day one, but a defensible posture at each — with depth added where criticality, provider concentration and jurisdictional exposure are greatest.

Connecting to the rest of the framework

Resilience & Recovery is the terminus, and it closes the loop. It consumes the criticality ratings from Visibility (G1), the residual risk that Protect and Oversight leave in place, and the assurance case that Assurance produces — turning those into continuity design and sovereignty posture proportionate to what each system actually is. When Oversight's kill switch and incident response (G15) prove insufficient, R&R is what preserves capability and control at the estate level.

Its outputs feed back into the beginning. Concentration risk, vendor viability findings, sovereignty exposure and continuity-exercise results all inform the next round of use-case approval, procurement decisions and system design — the recovery-and-improvement loop that turns the framework from a static structure into one that learns from what it has to handle.

About this series

This is Pillar 5 — the final pillar — of The 15 Guardrails of AI, published by Vertex11. The full series covers Visibility, Protect, Oversight, Assurance, and Resilience & Recovery, each aligned to the guardrails, controls, and standards mapping in the V11 AI Governance Framework. Together they set out a single, standards-validated approach to governing AI across its lifecycle — from the systems you have, through the perimeter around them and the oversight over them, to the evidence they meet the rules and the arrangements that keep them running when things fail.

To pressure-test the continuity and sovereignty posture around your AI, get in touch.