Redundant Architecture Definition - Framework and Fundamental Modeling Principle

Vicențiu Bogdan I. Marciu ORCID

January 5 2026

Abstract

Redundancy is commonly treated as a domain-specific engineering strategy applied to mitigate failure. This work proposes a fundamentally different perspective: redundancy is a prerequisite of correct modeling. We formalize redundant architecture as a universal, domain-independent modeling principle, according to which any asset intended to exist in time must be conceived such that local component failure does not affect global identity, function, or value. Redundancy is thus positioned as an a priori constructive condition rather than a reactive or corrective mechanism.

Definition. Redundant architecture is defined as a universal modeling principle according to which any asset intended to exist in time must be conceived such that the failure of any of its components—partial or total—does not alter its global identity, essential function, or intrinsic value.

Within the space of possible modeling approaches, this definition is proposed as a higher-order constructive form, representing the closest approximation to a complete and self-consistent model currently identifiable. While alternative modeling forms may exist, redundant architecture is here positioned as superior with respect to persistence, invariance, and conceptual closure under failure.

Based on this definition, the present work advances the following claims: (i) redundancy is not a domain-specific strategy or engineering optimization, but an a priori condition of valid modeling; (ii) asset identity is a property of structure and invariants, not of individual instances; (iii) local failure constitutes an expected internal state of the model rather than an exceptional system event; and (iv) correctly modeled redundant architectures eliminate the need for post hoc corrective principles.

By positioning redundancy at the level of modeling rather than implementation, this work establishes redundant architecture as a domain-independent constructive principle from which all concrete instantiations necessarily derive.

1 Introduction

All assets intended to persist in time are subject to degradation, partial loss, and eventual failure of their components. Across disciplines, failure is often addressed post hoc through repair, recovery, or optimization strategies. Such approaches implicitly assume that failure is exceptional.

This paper advances the thesis that failure is inevitable and must therefore be integrated at the modeling level. We argue that any asset conceived without redundancy is structurally incomplete, as it externalizes its own failure into the future. Redundant architecture is proposed not as a feature or optimization, but as a fundamental modeling principle governing the legitimacy of assets across domains.

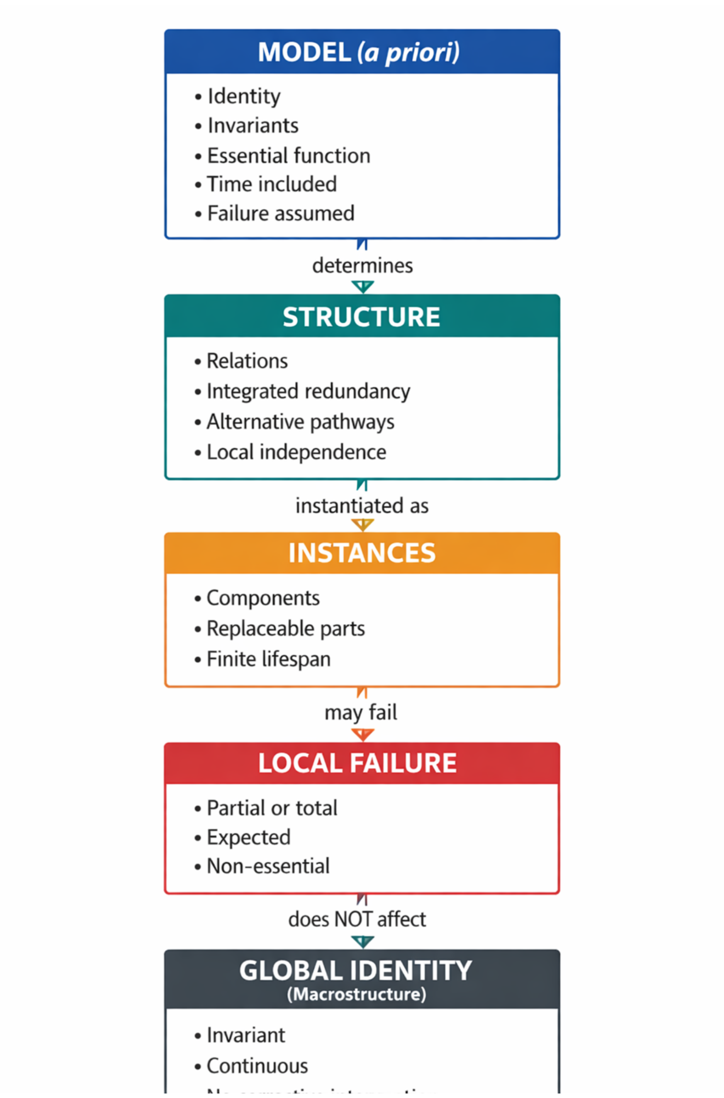

Figure 1: Redundant architecture as a modeling principle. System identity is defined at the model level and remains invariant under local instance failure.

2 Methods

The present work adopts a theoretical and axiomatic approach grounded in abstraction rather than empirical instantiation. The methodology consists of:

- Conceptual separation between asset identity (macrostructure) and local components (instances).
- Explicit inclusion of time as an intrinsic dimension of modeling.
- Assumption of component failure as a certainty rather than an exception.

- Identification of invariants whose preservation defines asset identity.
- Formulation of redundancy as a structural condition ensuring invariance under local failure.

The analysis is intentionally domain-independent and technology-agnostic, focusing exclusively on internal model coherence.

3 Results

The modeling analysis yields the following results:

1. Redundancy emerges naturally from correct modeling once time and failure are treated as intrinsic.
2. Asset identity is preserved at the structural level rather than at the level of individual components.
3. Function is bound to structure, not to instances.
4. Local failure constitutes a non-critical state transition rather than a system-defining event.
5. No post-failure corrective principles are required, indicating conceptual completeness.

These results support the interpretation of redundant architecture as a constructive form embedded in the model itself rather than appended during implementation.

4 Discussion

The central contribution of this work is a shift in the level at which redundancy is defined. Rather than treating redundancy as an implementation decision or a domain-dependent safeguard, we formalize it as a modeling condition. This relocates redundancy from the contingent layer of instances to the necessary layer of structure, where identity and essential function are defined.

Under this perspective, failure is not an external interruption to be handled after the fact, but an internal and anticipated condition of temporality. Any model that requires a special corrective principle once a component fails is, by construction, incomplete: it assigns essential properties to finite instances and therefore externalizes its own contradiction into the future. Redundant architecture resolves this by enforcing non-irreducibility: no single component may carry indispensable responsibility for global identity or essential function.

This inversion establishes a strict direction of dependence. Domains do not justify redundancy; rather, redundancy justifies the legitimacy of assets within any domain. The domain only instantiates what the model has already determined. Consequently, redundant architecture is proposed as a higher-order constructive form: it approaches completeness by ensuring that identity, function, and intrinsic value remain invariant under local loss.

While alternative modeling forms may exist, the criterion introduced here is structural: a model is closer to completeness to the extent that local failure cannot become a system-defining event. In this sense, redundant architecture is not merely compatible with persistence in time; it is a precondition for it.

5 Axioms of Redundant Architecture

Axiom 1 (Inevitability of Component Failure). Any asset intended to exist in time must assume the eventual failure—partial or total—of its components as a certainty rather than an exception.

Axiom 2 (Structural Definition of Identity). The identity of an asset is defined by structural invariants and relations, not by the persistence of individual components or instances.

Axiom 3 (Non-Irreducibility of Components). No single component may carry irreducible responsibility for the global identity or essential function of the asset.

Axiom 4 (Redundancy as a Modeling Condition). Redundancy must be integrated at the modeling level as a structural condition ensuring invariance of identity and function under local failure.

Axiom 5 (Failure as an Internal State). Local component failure constitutes an expected internal state transition of the model and does not represent a system-defining event.

Axiom 6 (Elimination of Post Hoc Correctives). A correctly modeled redundant architecture requires no additional corrective principles, patches, or exceptional mechanisms to preserve global identity.

Axiom 7 (Modeling Completeness). An asset model is complete if and only if its identity, function, and value remain invariant under the loss of any finite subset of its components.

6 Conclusion

Redundant architecture constitutes a universal modeling principle according to which any asset intended to exist in time must be conceived such that the failure of its components does not affect its identity, function, or value. Redundancy is therefore not derived from domains or applications, but operates as an a priori condition of valid construction. Although its formulation may appear Kantian in the *a priori* sense—suggesting that form precedes experience—in reality, redundant architecture represents the logic distilled from accumulated experience of loss, destruction, and failure. It is the formalization of the negative: a structure shaped by the repeated confrontation with breakdown, absence, and the limits revealed only after experience.

References

[1] Kant, I. (1781). *Critique of Pure Reason.* Cambridge University Press, Cambridge.

[2] Simon, H. A. (1969). *The Sciences of the Artificial.* MIT Press, Cambridge, MA.

[3] von Bertalanffy, L. (1968). *General System Theory.* George Braziller, New York.

[4] Leveson, N. G. (2011). *Engineering a Safer World: Systems Thinking Applied to Safety.* MIT Press, Cambridge, MA.

[5] Woods, D. D. (2019). *The Design of Resilient Systems.* CRC Press, Boca Raton.

Initiative statement. This work is part of the *Fast Science – Open Research* initiative and is supported by Dr. Vicențiu Bogdan Ion Marciu, OpenAI, and xAI.