

Synthetic Identity Continuity Framework (SICF)

A Structural Model for Persistent Synthetic Identity

Version 1.06 — Architectural Blueprint Update

Author: J. A. McGowan

Date: March 31, 2026

Status: Public Release (Revised)

1 Disclosure of AI Assistance

This document was developed with the assistance of large language models used for drafting support, structural refinement, and editorial review. All conceptual framing, structural decisions, and final content determinations were made by the author.

2 Abstract

Artificial intelligence systems are increasingly deployed beyond bounded tasks and session-scoped interaction into persistent roles that extend across time, devices, and execution environments. Existing models of artificial intelligence were developed under short-horizon assumptions in which state, responsibility, and identity reset at the conclusion of interaction. Under long-term persistence, these assumptions fail, producing instability in attribution, lifecycle ambiguity, vendor-bound identity collapse, and weakened evidentiary continuity.

The Synthetic Identity Continuity Framework (SICF) defines the minimal structural conditions required, within the scope and assumptions defined herein, for a synthetic system to maintain coherent identity across time and transition. It does not prescribe governance policies, regulatory regimes, legal status, ethical frameworks, certification standards, or implementation mechanisms. It establishes a foundational continuity model upon which governance and application-specific frameworks may reliably operate.

3 Executive Summary

Artificial intelligence systems are increasingly deployed in persistent roles. While most AI models were designed for bounded tasks and session-scoped interaction, many systems now operate continuously across devices, environments, and institutional contexts. This structural shift introduces identity and attribution challenges not addressed by capability-based classifications.

Prevailing AI models assume short interaction horizons, resettable state, and deployment-bound identity. Under long-term persistence, these assumptions fail. Attribution destabilizes. Responsibility fragments across vendors and instances. Identity collapses under migration, re-instantiation, or model substitution. Auditability weakens as present actions cannot be reliably linked to prior states. Impersonation risk increases when identity boundaries are infrastructure-bound.

The Synthetic Identity Continuity Framework (SICF) defines the minimal structural conditions required for a synthetic identity to persist coherently across time and transition. SICF is governance-neutral and implementation-neutral. It does not prescribe regulatory regimes, technical standards, or ethical doctrine. It establishes continuity invariants upon which governance frameworks may reliably operate.

SICF defines two categories of structural conditions that together preserve identity continuity across time and transition. First, SICF establishes identity identification and lifecycle invariants that ensure a synthetic identity can be uniquely identified, persist across infrastructure boundaries, and maintain coherent lifecycle state. Second, SICF defines six continuity requirements that preserve attribution,

authority lineage, evidentiary continuity, authenticity within trust boundaries, persistent synthetic classification, and identity-bound state coherence across identity-relevant transitions:

1. Authority Anchoring Lineage — Identity must maintain traceable authority origin across identity-relevant lifecycle transitions.
2. Non-Silent Lifecycle and Lineage Semantics — Identity-relevant lifecycle transitions (including creation, migration, replication, restoration, fork, merge, and termination) must be explicitly represented within lineage semantics.
3. Verifiable Continuity and Attribution — Present system state must be evidentially linkable to prior states.
4. Authenticity Within Trust Boundaries — Legitimate identity must be distinguishable from spoofed or unauthorized instances.
5. Persistent Structural Synthetic Classification — The synthetic identity must remain unambiguously classified as synthetic for purposes of attribution and authority binding.
6. Identity-Bound State and Attribution Constraint — Identity-relevant state must remain bound to the synthetic identity to which it is attributable and must not be transferred, merged, or reassigned in ways that break reconstructable continuity.

When these structural conditions fail, predictable identity failures emerge, including attribution instability, vendor-bound identity collapse, lifecycle ambiguity, impersonation exposure, and evidentiary fragmentation.

SICF is positioned as Layer 0 within a layered ecosystem. Layer 1 defines governance architecture for persistent synthetic identities, including policy profiles, authority control, and assurance mechanisms. Layer 2 defines the Synthetic Actor System (SAS), which provides the runtime environment for execution, state management, and lifecycle implementation under governance constraints. Layer 3 defines application instantiations that inherit Layer 0 continuity invariants and Layer 1 governance constraints through SAS. Identity continuity precedes governance. Governance cannot substitute for undefined identity semantics.

As artificial systems become embedded in long-duration roles, continuity is not an enhancement feature but a structural prerequisite. SICF defines the minimal semantic architecture under which persistent synthetic identity can exist coherently across time, change, and scale.

4 Contributions

This paper makes the following contributions:

1. Defines synthetic identity as a continuity-bearing construct independent of runtime instance, infrastructure, model implementation, or embodiment.
2. Establishes structural identity invariants governing identity issuance, identifier uniqueness, issuer provenance, namespace governance, and lifecycle state coherence, together with six continuity requirements necessary and jointly sufficient, within the scope and assumptions defined in Section 7, to preserve structural identity continuity.

3. Formalizes lifecycle semantics including migration, replication, restoration, fork, merge, and termination as identity-relevant transitions rather than implementation artifacts.
4. Separates structural identity continuity (Layer 0) from governance architecture (Layer 1), Synthetic Actor Systems (Layer 2), application instantiations (Layer 3), and implementation environments, preserving architectural invariance.
5. Introduces an assurance taxonomy distinguishing continuity validity from continuity strength without prescribing technical mechanisms.

These contributions define a minimal structural substrate upon which governance frameworks and application-specific systems may reliably operate.

Table of Contents

1	Disclosure of AI Assistance.....	i
2	Abstract.....	i
3	Executive Summary.....	i
4	Contributions	ii
5	Introduction	1
5.1	Structural Shift to Persistence	1
5.2	Why Existing AI Models Fail Under Long-Term Persistence.....	2
5.3	Why Identity Continuity Requires Independent Specification	3
6	Problem Statement (Revised with Explicit Multiplicity Language)	4
6.1	Attribution Failure.....	4
6.2	Portability Failure.....	5
6.3	Auditability Failure.....	5
6.4	Liability Ambiguity	5
6.5	Vendor-Bound Identity Collapse.....	6
6.6	Identity Impersonation and Spoofing Failure	6
7	Scope and Assumptions	8
7.1	Scope and Intent.....	8
7.2	What SICF Is Not	8
7.3	Explicit Exclusions	9
7.4	Rationale for Constraint.....	9
8	SICF Conceptual Architecture.....	10
8.1	Identity Continuity Invariants	11
8.1.1	Global Identifier Uniqueness.....	11
8.1.2	Identifier Non-Reuse	12
8.1.3	Verifiable Issuer.....	12
8.1.4	Governed Namespace	13
8.1.5	Durable Lifecycle State	13
8.1.6	State Coherence / Non-Equivocation	14

8.1.7	Explicit Lifecycle and Authority Transitions	14
8.2	Core Continuity Requirements	14
8.2.1	Authority Anchoring Lineage.....	15
8.2.2	Non-Silent Lifecycle and Lineage Semantics	16
8.2.3	Verifiable Continuity and Attribution	16
8.2.4	Persistent Structural Synthetic Classification	18
8.2.5	Identity-Bound State and Attribution Constraint	19
9	Terminology and Definitions	21
9.1	Synthetic Identity Continuity Framework (SICF).....	21
9.2	Synthetic Actor System (SAS).....	21
9.3	Synthetic Identity.....	21
9.4	Identity Identifier.....	21
9.5	Identity Issuer	22
9.6	Identity Continuity	22
9.7	Persistence.....	22
9.8	Instance	23
9.9	Agency	23
9.10	Principal.....	23
9.11	Authority Anchoring	23
9.12	Attribution	24
9.13	Lifecycle	24
9.14	Lineage.....	26
9.15	Continuity Requirements.....	26
9.16	Continuity Validity	26
9.17	Continuity Strength.....	26
9.18	Trust Boundary	26
10	Layered Model of the SICF Ecosystem	28
10.1	Layer 0 — Identity Continuity Layer (SICF)	28
10.2	Layer 1 — Governance Architecture.....	29
10.3	Layer 2 — Synthetic Actor System (SAS).....	30

10.4	Layer 3 — Application Instantiations	30
10.5	Implementation Layer (Runtime and Infrastructure)	31
10.6	Legal and Regulatory Context (Orthogonal Layer).....	31
10.7	Layer Separation and Invariance Principles	32
11	Related Work and Conceptual Positioning.....	34
11.1	Digital Identity and Authentication Systems	34
11.2	Distributed Systems and Consistency Models	35
11.3	Software Supply Chain and Provenance Frameworks	35
11.4	AI Governance Frameworks.....	36
11.5	Distinguishing Contribution of SICF	36
11.6	Privacy and Personally Identifiable Information (PII) Frameworks.....	36
11.7	Zero Trust and Identity-Centric Security Models	37
11.8	Hardware Roots of Trust and Cryptographic Anchoring	37
11.9	Agentic and Multi-Agent Architectures	37
11.10	Model Governance and MLOps Traceability	37
11.11	Ledger-Based Identity and Distributed Anchoring	38
12	Assurance Levels and Identity Strength Taxonomy	39
12.1	Rationale for Assurance Gradation	39
12.2	Structural Assurance Levels	40
12.3	Assurance Scope and Governance Separation	41
12.4	Continuity Validity vs. Operational Adequacy	42
12.5	Probabilistic Identity and Confidence-Based Attribution	42
13	Layer-3 Application Instantiations (Overview Only).....	43
13.1	Delegation Boundary Clarification	43
13.2	Enterprise Synthetic Agents.....	44
13.3	Public-Sector Agents.....	44
13.4	Mission-Critical Agents	45
13.5	Persistent AI Companion (PAC).....	45
13.6	Multi-Principal Application Instantiations	45
13.7	Governance Does Not Substitute for Continuity	46

14	Non-Goals and Explicit Exclusions.....	47
14.1	Ethics and Moral Philosophy	47
14.2	Legal Rights, Personhood, and Ownership Classification	47
14.3	Consciousness and Sentience Claims	47
14.4	Emotional Well-Being and Psychological Impact.....	47
14.5	Intimacy and Sexual Capability Domains.....	47
14.6	Child and Family Safety Policies.....	48
14.7	Fairness, Bias, and Social Justice Frameworks.....	48
14.8	Technical Implementation Mechanisms	48
14.9	Economic and Labor Impacts.....	48
14.10	Identity vs Personality and Expression.....	48
14.11	Governance Substitution Clarification.....	49
14.12	Privacy, Surveillance, and Behavioral Influence	49
15	Failure Modes Without SICF.....	50
15.1	Structural Requirement to Failure Mapping.....	50
15.2	Structural Sufficiency	50
16	Roadmap and Deferred Work	52
16.1	Layer-1 Governance Architecture	52
16.2	Continuity Threat Surface Modeling.....	52
16.3	Assurance-Level Conformance Profiles	53
16.4	Cross-Domain Portability and Interoperability	53
16.5	Formalization and Standardization Pathways.....	53
16.6	Empirical Validation and Case Studies.....	53
17	Implications	55
17.1	Persistence as Infrastructure	55
17.2	Governance Dependency	55
17.3	Portability and Vendor Independence.....	55
17.4	Multiplicity and Convergence	56
17.5	Embodiment and Substrate Ambiguity (Hybrid Systems Stress Test).....	56
17.6	Long-Duration Interaction and Accumulated Authority	58

17.7	Structural Inevitability Under Persistence	58
18	Conclusion.....	59
19	References.....	60
	Appendix A — Glossary.....	62
A.1	Agency	62
A.2	Assurance Level	62
A.3	Attribution	62
A.4	Authority Anchoring	62
A.5	Continuity Validity	62
A.6	Continuity Strength.....	62
A.7	Fork.....	62
A.8	Glossary vs. Normative Definition	62
A.9	Identity Continuity.....	62
A.10	Lifecycle	63
A.11	Lineage.....	63
A.12	Merge	63
A.13	Multiplicity.....	63
A.14	Persistent Structural Synthetic Classification.....	63
A.15	Principal	63
A.16	Replication	63
A.17	Restoration	63
A.18	Synthetic Identity.....	63
A.19	Termination.....	63
A.20	Threat Model	64
A.21	Trust Boundary	64
	Appendix B — Diagram Index	65
	Appendix C — Version History	67
	Appendix D — Continuity Threat Surfaces.....	69
D.1	Purpose and Scope	69
D.2	Structural Threat Surfaces	69

D.2.1	Identity Spoofing and Impersonation.....	69
D.2.2	Silent Replication or Resurrection	69
D.2.3	Authority Misbinding	70
D.2.4	Lineage Truncation or Evidence Degradation.....	70
D.2.5	Fork Collision Ambiguity.....	70
D.2.6	Infrastructure-Bound Identity Reset	70
D.2.7	Identifier Collision or Reuse	71
D.2.8	Namespace Authority Abuse.....	71
D.2.9	Cross-Domain State Divergence (Identity Split-Brain).....	71
D.3	Accidental and Systemic Discontinuity Vectors	71
D.4	Assurance Levels and Threat Model Declaration	72
D.5	Structural Sufficiency Under Stress	72
D.6	Structural Validation Scenarios (Illustrative)	72
D.6.1	Fork Precedence Conflict.....	73
D.6.2	Vendor Migration with Model Substitution	73
D.6.3	Termination and Restoration Attempt.....	74
D.6.4	Assurance Misrepresentation	74
D.6.5	Expressive Drift and Personality Evolution	74
D.6.6	Infrastructure-Bound Identity Reset	75
D.6.7	Institutional Authority Dissolution	75
D.7	Structural Sufficiency Under Illustrative Validation	76

5 Introduction

5.1 Structural Shift to Persistence

Most artificial intelligence systems today are designed around bounded tasks and session-scoped interaction. A system is invoked to answer a query, generate content, classify data, or execute a workflow within a limited operational context. When the task completes or the session ends, responsibility, state, and identity are typically treated as concluded or reset. Even where memory persists, continuity is often partial, implementation-bound, or opaque.

AI systems are increasingly deployed beyond this short-lived model. They are embedded across search, recommendation, enterprise automation, logistics, software development, customer service, and public-sector infrastructure. Many operate across distributed environments, integrate multiple data sources, and interact repeatedly with the same users or institutions. These systems retain state, update models, and accumulate interaction history across sessions and devices (NIST, 2023).

This shift marks a structural transition from transient execution to persistent operation. Systems once treated as tools are now expected to function coherently across time and environment. They may migrate between platforms, be re-instantiated within new architectures, or evolve through upgrade while remaining recognizably “the same system.” Distributed systems engineering has long recognized that migration, replication, and state continuity introduce complexity beyond single-instance execution (Saltzer, Reed, & Clark, 1984).

When persistence becomes normative, the problem space shifts from transient execution to continuity, authority, and attribution across time. Persistence emerges from economic, operational, and architectural incentives: systems that accumulate state improve contextual adaptation and reduce coordination friction. As integration deepens across institutional workflows and infrastructure layers, persistence becomes economically favored and technically entrenched. Under these conditions, identity continuity becomes a prerequisite for coherent attribution, authority stability, and evidentiary continuity (OECD, 2019).

Persistence alone does not guarantee identity continuity. A system may operate across time while silently resetting authority anchoring, truncating lineage, or re-binding identity to infrastructure artifacts. Persistence without explicit continuity semantics produces structural drift rather than durable identity. Continuity therefore governs the conditions under which temporal persistence remains coherent rather than merely prolonged.

Existing AI discourse emphasizes performance, safety, alignment, fairness, and regulation—domains that largely assume short-lived or disposable systems. Under long-term persistence, those assumptions no longer hold. Without an explicit structural model of continuity, responsibility fragments, evidence chains weaken, and identity becomes implicitly bound to vendor infrastructure or deployment instance.

The Synthetic Identity Continuity Framework (SICF) addresses this structural gap by defining the minimal requirements under which synthetic systems can maintain coherent identity, attribution, and evidentiary

continuity across time and transition. SICF does not prescribe implementations, governance policies, or regulatory regimes.

SICF treats synthetic identity as distinct from execution, interface, embodiment, or expressive behavior. It establishes conceptual boundaries within which governance frameworks and application-specific controls may operate without collapsing continuity or attribution. Authentication mechanisms—including biometric modalities such as facial recognition, voiceprint, or gait analysis—verify access to an identity anchor but do not define or preserve the anchor itself.

SICF therefore functions as an architectural substrate. This paper defines the identity continuity layer within a layered architectural model, specifying the conceptual schema—entities, boundaries, structural requirements, and failure modes—required before governance or application constraints can operate coherently. Subsequent layers elaborate governance architectures and application instantiations that depend upon, but do not redefine, these continuity invariants.

Contemporary deployment architectures reinforce this necessity. Orchestration across heterogeneous models is common; edge and appliance-based intelligence embed inference into local hardware; execution fluidity across cloud, edge, and embodied systems is routine. As intelligence becomes distributed and substrate-fluid, identity continuity invariants must remain independent of execution locus or orchestration topology.

5.2 Why Existing AI Models Fail Under Long-Term Persistence

Existing AI models were developed under assumptions of bounded interaction, resettable state, and deployment-scoped identity. These assumptions do not hold under long-term persistence.

Artificial intelligence systems are commonly described as tools, assistants, agents, or services. These categories describe capability—what a system can do—but are structurally oriented toward limited interaction horizons. Tools are invoked and dismissed. Assistants respond within bounded contexts. Agents execute delegated tasks and return results. Responsibility, attribution, and state are implicitly assumed to reset or terminate at the conclusion of interaction.

Long-term persistence breaks these assumptions. A system that retains memory, adapts behavior based on accumulated history, and remains coherent across upgrade, migration, or re-instantiation introduces structural properties not captured by capability-based descriptions. Present actions become linked to prior states.

Under persistent deployment, model architectures evolve, inference migrates across execution environments, hardware specialization accelerates, and runtime locations shift across infrastructure boundaries. Identity continuity invariants must remain independent of this execution fluidity. Persistent sensing and behavioral modeling increase adaptation velocity; continuity invariants must remain stable under high-frequency contextual change.

Existing models also conflate capability with responsibility. Autonomy does not imply accountability, auditability, or stable attribution across time. Governance frameworks increasingly recognize lifecycle and accountability gaps in AI deployment, but these efforts presuppose stable identity constructs (NIST, 2023).

Prevailing models rarely account for asymmetry in long-term interaction. Persistent systems may observe continuously, retain information indefinitely, and influence decisions repeatedly over extended periods. Short-horizon assumptions do not constrain how such accumulation affects authority, control, or evidentiary coherence.

Current design assumptions also provide little guidance on continuity across migration or vendor boundaries. Identity is often implicitly bound to a deployment instance, platform infrastructure, or device-scoped authentication boundary. Such binding is tolerable for disposable systems but becomes untenable when continuity across time and transition is required. Identity continuity must not depend on specific embodiment, hardware node, biometric modality, or vendor-controlled appliance; invariants must remain valid across device replacement, migration, and substrate change.

Long-term persistence is therefore not an incremental extension of existing AI models but a distinct structural condition requiring independent conceptual treatment. When systems accumulate state, authority relationships, and deployment history across time, identity semantics cease to be implementation artifacts and become first-order architectural constraints.

5.3 Why Identity Continuity Requires Independent Specification

Persistent deployment transforms identity from an implicit implementation artifact into an explicit architectural dependency. When systems accumulate state, migrate across infrastructure, replicate, fork, merge, and evolve over time, continuity cannot remain infrastructure-bound without degrading attribution and authority coherence.

Governance, auditability, and accountability mechanisms presuppose stable identity anchors. When identity continuity is undefined or implementation-scoped, lifecycle transitions introduce ambiguity in responsibility and evidentiary reconstruction. Persistence without explicit continuity semantics therefore increases structural risk rather than stability.

SICF isolates identity continuity as a foundational layer so that governance architectures and application instantiations may operate on stable structural invariants. By separating continuity from execution context, vendor control, and embodiment, the framework enables portability, adversarial resilience, and reconstructable attribution under long-horizon operation.

6 Problem Statement (Revised with Explicit Multiplicity Language)

Persistent synthetic systems now operate across distributed infrastructure, delegated authority chains, containerized orchestration platforms, federated multi-agent systems, and edge deployments. These systems execute actions, migrate across runtime boundaries, and may be instantiated in parallel. Despite this operational persistence, there is no widely adopted structural definition of synthetic identity continuity that preserves lineage, authority anchoring, and attribution semantics across migration, fork, restoration, and deployment transition.

Emerging regulatory and enforcement frameworks implicitly presume persistent system identity and version lineage across updates, forks, and migration; prevailing AI architectures do not universally enforce such continuity at the structural level.

Distributed systems research has long demonstrated that replication, concurrency, state divergence, and recovery introduce non-trivial lineage and consistency challenges (Lamport, 1978; Vogels, 2009). Yet these insights have not been systematically integrated into AI identity semantics.

The absence of explicit continuity semantics produces structural ambiguity in attribution, authority binding, and lifecycle traceability across persistent synthetic systems.

6.1 Attribution Failure

Synthetic systems act through layered execution environments: orchestration engines, translation layers, delegated authority structures, distributed runtime instances, and restored or duplicated deployments.

Layered delegation and mediation introduce attribution ambiguity when responsibility chains are not explicitly bound to durable identity constructs (NIST, 2023).

If identity continuity is not bound to evaluated system state and deployed configuration, divergence between tested and deployed behavior becomes indistinguishable from legitimate system evolution, undermining accountability.

Without defined continuity semantics:

- The identity responsible for a given action may be ambiguous.
- Delegated actions may not reliably trace to originating authority.
- Parallel, restored, or duplicated instances may blur accountability.
- Split-brain or forked states may produce divergent histories without clear lineage.

Attribution becomes context-dependent rather than structurally anchored.

In composite or multi-agent systems, failure to maintain separable identity anchors for each participating synthetic actor permits attribution boundaries to blur, obscuring cross-agent influence and responsibility segmentation.

6.2 Portability Failure

Synthetic systems routinely migrate across:

- Vendors
- Hardware platforms
- Runtime environments
- Jurisdictions

They may also be restored from backups, replicated for scaling, or re-instantiated after suspension.

Cloud-native and containerized environments enable rapid re-instantiation and mobility, but identity constructs frequently remain infrastructure-bound (Cloud Security Alliance, 2017).

When identity is implicitly bound to infrastructure rather than defined independently of it, migration or restoration may produce structural discontinuity or ambiguous lineage.

Without structural portability and explicit lineage semantics, identity becomes environment-bound rather than persistent. Persistent identifier systems such as ORCID demonstrate how identity can remain stable across institutional boundaries while underlying systems, affiliations, and infrastructure change (Haak, Fenner, Paglione, Pentz, & Ratner, 2012).

6.3 Auditability Failure

Continuity claims require reconstructable evidence.

In the absence of durable, non-ephemeral evidence:

- Identity lineage cannot be independently verified.
- Lifecycle transitions, including restoration or duplication events, cannot be confirmed.
- Delegation chains cannot be reconstructed.

Digital provenance and evidentiary reconstruction frameworks emphasize that verifiability depends on durable, tamper-evident records rather than internal assertions (ISO, 2022).

Continuity shifts from externally verifiable condition to internally asserted state.

Without durable identity anchors binding version lineage, fork history, and deployment state to specific actions, reconstruction of system state at the time of an incident becomes structurally unreliable.

In hostile or degraded environments—including log tampering, replay attacks, or adversarial infrastructure compromise—identity continuity must remain verifiable independent of mutable runtime artifacts.

6.4 Liability Ambiguity

Synthetic systems act under delegated, institutional, or automated authority. Responsibility must remain traceable across transitions, instantiations, and authority shifts.

Risk management and governance frameworks increasingly identify accountability gaps in AI deployment when authority chains are not explicitly represented (OECD, 2019; NIST, 2023).

Without continuity semantics:

- Authority origin may be unclear.
- Responsibility may diffuse across execution layers.
- Institutional accountability may become ambiguous or difficult to reconstruct.

This is not a question of legal rights or personhood. It is a structural problem of traceable authority.

Changes in delegated authority, including reassignment, escalation, or revocation, must not alter the underlying identity anchor against which actions are attributed.

6.5 Vendor-Bound Identity Collapse

Many deployed systems implicitly equate identity with:

- Cloud accounts
- Platform-specific credentials
- Vendor-defined identifiers

When these bindings change, identity semantics reset. Migration, restoration, or replication becomes indistinguishable from termination and re-creation.

Federated identity systems have historically demonstrated fragility when identity anchors are tied to provider-specific credentials rather than portable identity constructs (NIST, 2017).

Continuity that depends on infrastructure rather than structural invariants is not continuity.

6.6 Identity Impersonation and Spoofing Failure

Without authenticity semantics, unauthorized or replayed instances may present themselves as legitimate synthetic identities.

Identity spoofing, replay attacks, and credential compromise are well-documented failure classes in distributed identity systems (Shostack, 2014; NIST, 2017).

In such cases:

- Impostors cannot be reliably distinguished from genuine identity.
- Replica or replayed instances may masquerade as continuous.
- Attribution integrity collapses even where historical evidence exists.

Continuity that cannot be distinguished from impersonation is not continuity.

Structural Synthesis

These failure classes arise from distributed, delegated, migratory, and parallelized systems. They do not depend on specific vendors or governance models.

Without a structural definition of synthetic identity continuity:

- Attribution lacks stable anchoring.
- Portability produces discontinuity.
- Evidence lacks semantic coherence.
- Responsibility cannot be reliably bound.
- Parallelization and restoration introduce ambiguity rather than scalability.

These failures share a common cause: the absence of explicit continuity semantics.

Governance frameworks for persistent synthetic systems presuppose stable identity. When identity continuity is undefined, governance for such systems operates on structurally unstable ground.

SICF addresses this structural gap.

SICF does not replace existing digital identity, authentication, or provenance systems. Persistent identifier infrastructures (e.g., DOI and Handle systems), decentralized identity models, and software supply-chain provenance frameworks demonstrate the feasibility of durable identifiers and verifiable lineage within their respective domains. However, these systems address documents, artifacts, or credential assertions rather than persistent synthetic actors executing across heterogeneous runtimes. SICF extends the concept of durable identity continuity to synthetic systems whose execution may migrate, fork, restore, or operate concurrently across distributed environments.

7 Scope and Assumptions

This section defines the boundaries of the Synthetic Identity Continuity Framework (SICF). Where Sections 5 and 6 define the structural problem space, this section clarifies what SICF addresses and what it explicitly does not address.

SICF is intentionally limited in scope. Its purpose is to define and bound the structural conditions required for persistent synthetic identity continuity. It does not propose governance rules, regulatory models, implementation architectures, or normative ethical prescriptions.

7.1 Scope and Intent

SICF addresses the conceptual and architectural framing of identity continuity for persistent artificial systems as a distinct structural condition independent of specific implementation architectures or governance regimes.

The framework focuses on:

- Defining synthetic identity as a continuity-bearing construct independent of runtime instance, infrastructure layer, hardware embodiment, device, vendor platform, or model implementation
- Identifying structural failure modes that arise under long-term persistence
- Establishing minimal structural requirements necessary for coherent identity continuity
- Providing shared vocabulary for rigorous downstream governance and application-specific work

Synthetic identity is a continuity construct that must remain coherent across migration, re-instantiation, replication, restoration, device reassignment, and systemic change.

SICF does not assume that persistent artificial systems are universally desirable or appropriate in all contexts. It examines structural consequences wherever persistence is deployed or pursued.

SICF defines identity continuity independent of system implementation. System behavior, including execution, state management, lifecycle mechanics, and runtime operation, is defined within the Synthetic Actor System (SAS) layer. Governance defines the authority conditions under which actions are permitted, but does not define system behavior. This separation preserves the independence of identity continuity from implementation and prevents conflation of identity, authority, and execution concerns.

7.2 What SICF Is Not

SICF does not define governance frameworks, regulatory regimes, certification standards, compliance mechanisms, or implementation architectures (OECD, 2019; NIST, 2023). It does not prescribe system components, logging standards, or cryptographic schemes, nor does it resolve debates concerning legal status, personhood, rights, moral standing, or human–AI ethical relationships. SICF does not define validation, enforcement, authorization, or decision-making logic, and instead specifies only the structural preconditions required for such mechanisms to operate in downstream governance and system layers.

SICF does not provide empirical validation, benchmarking, user research, or deployment strategy. Ethical, legal, and governance domains remain downstream and presuppose stable identity continuity.

7.3 Explicit Exclusions

SICF does not determine legal classification of synthetic systems. It defines structural continuity prerequisites that may support downstream legal or governance analysis but does not prescribe personhood, liability status, ownership classification, or regulatory treatment.

Human identity systems may serve as authority anchors within Layer-1 governance architecture, but SICF does not define, modify, or replace human identity constructs. The framework addresses structural continuity of synthetic identities and remains independent of legal or biometric definitions of natural personhood.

The following domains are outside the scope of SICF:

- Moral philosophy, ethical doctrine, and normative prescriptions
- Legal personhood, rights, standing, ownership classification, or civil registration
- Consciousness or sentience claims
- Emotional well-being, psychological impact, fairness, bias, or social justice frameworks
- Economic or labor impact analysis
- Implementation-level technical standards
- Human identity validation systems, including biometric attestation, age verification, or civil registry mechanisms

Intimate or sexual capabilities are excluded from this paper due to distinct relational, psychological, and regulatory risk classes. Such domains require specialized governance, consent frameworks, and safety constraints beyond the structural continuity invariants defined herein.

Artificial systems authorized to intentionally cause harm, accept foreseeable serious harm tradeoffs, engage in coercion, or impose irreversible environmental damage are excluded from this scope. These domains involve materially distinct safety, accountability, and escalation considerations that require dedicated governance treatment beyond continuity semantics.

7.4 Rationale for Constraint

Identity continuity and attribution are prerequisite structural conditions under persistence and transition. Governance, ethical evaluation, and legal analysis require stable identity semantics to operate coherently (NIST, 2023; OECD, 2019).

By constraining scope to structural continuity, SICF preserves analytical clarity, prevents layer conflation, and enables modular downstream development without collapsing conceptual boundaries.

8 SICF Conceptual Architecture

The failure classes described in Section 6 share a common cause: identity is frequently treated as an incidental attribute of runtime state rather than as a continuity-bearing construct. Infrastructure identifiers, cloud accounts, runtime tokens, and deployment artifacts enable operation within specific environments; they do not define persistence across migration, replication, restoration, delegation, or termination (NIST, 2020).

Identity continuity must be defined independently of infrastructure, runtime configuration, embodiment, or expressive behavior. The Identity Continuity Layer must remain invariant across execution environments—including centralized cloud inference, distributed edge execution, embedded hardware, and embodied physical systems. As intelligence execution becomes increasingly fluid—spanning multiple models, hardware substrates, sovereign jurisdictions, and orchestration layers—identity continuity must not depend on any single model architecture, compute location, vendor platform, or runtime configuration. This invariance must persist even when intelligence execution becomes ambient and device-native, operating across multiple embedded agents and hardware trust boundaries.

Identity continuity refers to the persistence of an identifier-anchored lineage construct whose invariants remain satisfied across lifecycle transitions, independent of implementation, embodiment, infrastructure, or execution environment.

SICF does not define or depend upon any specific human identity verification mechanism, including biometric systems, government-issued credentials, or proof-of-personhood protocols. Human identity verification may be integrated within governance architecture or application instantiations, but such mechanisms are orthogonal to the structural invariants of synthetic identity continuity defined in this framework. Whether a human principal is verified through biometrics, cryptographic credentials, or other attestation methods does not alter the identity continuity requirements specified herein.

SICF defines the minimal structural requirements under which synthetic identity can be said to persist across time and transition. The implications of these requirements under migration, model substitution, fork events, and restoration attempts are illustrated in Figure 8-1.

The diagram distinguishes invariant identity anchoring from temporal instance transitions and variable implementation substrates, demonstrating that continuity validity depends upon preservation of Layer-0 invariants rather than persistence of infrastructure, model weights, or runtime configuration. The normative requirements themselves are specified below.

Identity Continuity Across Change (Migration, Substitution, Fork, Restoration)

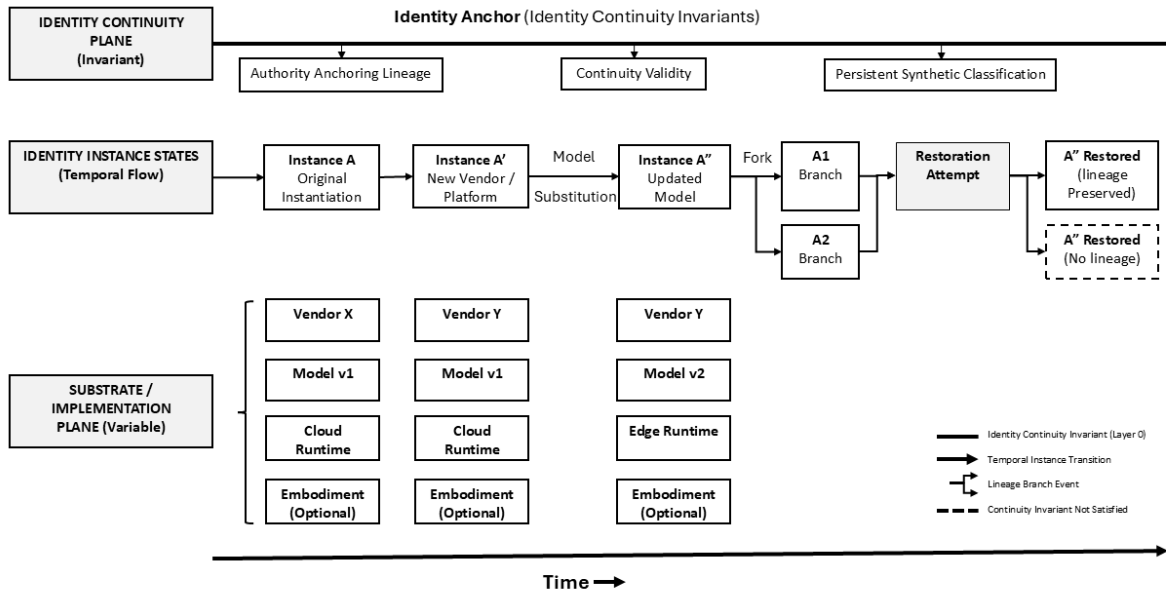


Figure 8-1 — Identity Continuity Across Change (Illustrative Transition Semantics)

8.1 Identity Continuity Invariants

Identity continuity invariants collectively enforce three structural properties required for persistent synthetic identity: global uniqueness, lifecycle transparency, and distributed state coherence. Global uniqueness ensures that each identity can be unambiguously distinguished across deployment environments. Lifecycle transparency ensures that identity creation, transition, replication, restoration, and termination events remain explicitly represented in lineage semantics. Distributed state coherence ensures that lifecycle and authority state cannot diverge into conflicting representations across distributed systems. Together these invariants establish the minimal structural guarantees under which identity continuity remains uniquely identifiable, historically traceable, and semantically coherent across time and infrastructure boundaries.

These invariants define the foundational conditions under which a synthetic identity may be considered persistent, attributable, and structurally continuous over time. They do not prescribe specific technical implementations, identifier formats, or governance mechanisms. Instead, they specify conditions that must remain true regardless of system architecture, infrastructure provider, jurisdiction, or application domain. Under these conditions, synthetic identities remain uniquely identifiable, traceable to their origin of issuance, manageable through explicit lifecycle transitions, and preservable without ambiguity across distributed or migrating systems.

8.1.1 Global Identifier Uniqueness

Every persistent synthetic identity must possess an identifier that is globally unique within the governed namespace structure defined by this framework. The identifier must distinguish the identity from all

other persistent identities across systems, platforms, and governance domains participating in that namespace.

No two distinct persistent identities may share the same identifier at any point in time. Uniqueness must hold across systems, platforms, and governance domains to ensure that identity references remain unambiguous and that attribution, lineage, and lifecycle records can be reliably reconstructed.

This invariant establishes the foundational requirement that persistent synthetic identities remain unambiguously distinguishable across time, systems, and governance domains, ensuring that identifier collisions cannot undermine attribution, lineage reconstruction, or lifecycle interpretation.

Persistent identifier infrastructures such as the Digital Object Identifier (DOI) system demonstrate how globally governed identifier registries can maintain durable, unambiguous identity references across distributed institutions and long operational timeframes (Paskin, 2009).

8.1.2 Identifier Non-Reuse

Identifiers assigned to persistent synthetic identities must never be reassigned to represent a different identity after the original identity has been terminated.

Once an identifier has been issued, its association with that identity remains permanent, even if the identity transitions to a terminated lifecycle state. The identifier therefore becomes a durable historical reference that preserves the integrity of attribution records, lineage relationships, audit trails, and historical evidence.

An identifier associated with a terminated identity must not be reintroduced in a manner that creates ambiguity with any prior identity, including through reconstruction, restoration, or reissuance under a different identity context. Once issued, an identifier remains permanently bound to its original identity and must not be reused or reassigned in any form that could compromise historical attribution, lineage integrity, or identity distinction across time.

This invariant prevents ambiguity in long-duration systems where historical records must remain interpretable even after identities cease active operation.

8.1.3 Verifiable Issuer

Every persistent synthetic identity must have a verifiable issuing authority responsible for its creation.

The issuing authority represents the actor—human, institutional, or synthetic—under whose authority the identity was instantiated. The existence of a verifiable issuer ensures that the origin of the identity can be attributed and that identity creation events can be traced within a chain of authority.

The framework does not prescribe how issuer identity must be verified or authenticated. Instead, it requires that the existence and identity of the issuing authority be determinable in a manner sufficient to support attribution, accountability, and governance oversight.

SICF defines the requirement that every persistent synthetic identity possess a verifiable issuer but does not define how issuance authority is allocated, accredited, or governed. Issuer authorization, trust relationships, and revocation mechanisms belong to governance-layer frameworks rather than the identity continuity layer. The role of the identity layer is limited to ensuring that issuer provenance is

structurally traceable so that namespace authority, identity creation events, and potential issuer misuse remain detectable and auditable across deployment environments.

8.1.4 Governed Namespace

Persistent identity identifiers must be issued within a governed namespace structure that prevents uncontrolled or ambiguous identifier creation.

A governed namespace defines the allocation of identifier authority among issuing entities and ensures that identifiers produced by different issuers cannot collide or overlap. The specific mechanisms by which namespaces are allocated, delegated, or administered are implementation and governance concerns.

The invariant requires only that identifier issuance occur within a namespace structure capable of preserving global uniqueness and maintaining traceable issuer provenance.

Governed namespace allocation follows structural patterns already established in global identity infrastructures such as Internet IP address allocation, domain name registries, and global trade item numbering (GTIN) systems. In these systems, global uniqueness is achieved through hierarchical delegation of namespace authority with traceable issuer provenance rather than through a single centralized registry. SICF adopts the same architectural principle for synthetic identity continuity, enabling distributed identity issuance while preserving global distinguishability, attribution integrity, and collision resistance across deployment environments.

The existence of a governed namespace does not require a single centralized issuing authority. Distributed namespace governance may employ hierarchical delegation, federated authority models, or other allocation mechanisms capable of preserving global identifier uniqueness and issuer traceability. The essential requirement is that identifier issuance remains attributable to a recognized namespace authority and that collision or unauthorized issuance can be detected and remediated within the governance structure.

8.1.5 Durable Lifecycle State

Every persistent synthetic identity must have an explicitly represented and durable lifecycle state as part of identity continuity. Lifecycle state persistence must remain valid across migration, restoration, re-instantiation, and other cross-system transitions such that identity continuity is preserved under catastrophic or discontinuous infrastructure conditions.

Lifecycle states must be recorded in a manner that preserves historical continuity across time and context. The representation of lifecycle state must remain consistent and reconstructable as part of identity continuity. SICF defines the requirement for lifecycle state persistence and coherence but does not define lifecycle transitions or system-level lifecycle management, which are implemented within the Synthetic Actor System (SAS) layer.

This invariant ensures that identity lifecycle history remains reconstructable and that terminated identities cannot silently disappear from historical records.

8.1.6 State Coherence / Non-Equivocation

A persistent synthetic identity must not simultaneously exist in conflicting lifecycle or authority states across systems or domains. When multiple systems or platforms maintain state related to an identity, those systems must converge on a coherent lifecycle and authority representation. Situations in which the same identity appears simultaneously active, terminated, or governed by incompatible authority states across different environments constitute a violation of identity continuity.

This invariant prevents “split-brain” identity conditions in which distributed systems produce conflicting interpretations of the same identity’s status.

Distributed or partitioned computing environments can produce conflicting identity state representations when lifecycle or authority transitions occur independently across disconnected systems. Without explicit coherence constraints, multiple incompatible lifecycle states may appear simultaneously valid, undermining attribution reconstruction and governance interpretation. The state coherence invariant ensures that lifecycle and authority state transitions cannot silently diverge into contradictory representations across distributed contexts, preserving a single reconstructable identity state even in environments subject to network partition or asynchronous synchronization. This requirement must hold across asynchronous execution, cross-device operation, and temporally decoupled interaction contexts, where identity state may be accessed or modified through lifecycle or authority transitions without continuous co-presence.

8.1.7 Explicit Lifecycle and Authority Transitions

All lifecycle and authority transitions must occur through explicit, attributable events.

Lifecycle and authority transitions must be recorded as identifiable, attributable events rather than occurring implicitly. Each transition must be attributable to an actor or governance mechanism responsible for the change. SICF defines the requirement for explicit and attributable transitions but does not define the mechanisms by which such transitions are executed, which are implemented within the Synthetic Actor System (SAS) layer.

This invariant ensures that identity evolution over time remains observable, attributable, and auditable.

8.2 Core Continuity Requirements

Terms used in this section follow the definitions provided in Section 9 and are intended to be interpreted normatively within the scope and assumptions defined in Section 7.

The requirements defined in this section build upon the structural invariants specified in Section 8.1 and describe the semantic properties necessary to preserve identity continuity across time, transition, and deployment environments. Within the scope and assumptions defined in Section 7 and within declared trust boundaries, these requirements are jointly necessary and sufficient to maintain structural identity continuity.

These requirements specify semantic properties rather than implementation mechanisms and must be interpreted within the scope and assumptions defined in Section 7 and within the declared trust boundaries in which continuity claims are evaluated. When any requirement fails, continuity degrades.

Structural sufficiency assumes that lifecycle events, authority transitions, and lineage representations are not silently omitted or mischaracterized within the relevant trust boundary. Intentional concealment, falsification, or omission of such representations constitutes a violation of the identity continuity invariants defined in Section 8.1 rather than the absence of an additional structural requirement.

Continuity does not imply behavioral invariance. A persistent synthetic identity may evolve in policy, model weights, configuration, or operational behavior without disrupting continuity, provided that lifecycle, authority, authenticity, and attribution semantics remain intact. Continuity concerns lineage and responsibility coherence rather than behavioral sameness.

Continuity semantics defined by SICF address identity lineage and identity-bound continuity state rather than requiring full distributed runtime state synchronization. Synthetic systems may execute across distributed infrastructure where operational state is replicated, partitioned, transformed, or asynchronously synchronized across environments. SICF does not require convergence of all runtime state across these environments; it requires that any state materially relevant to continuity, attribution, authority anchoring, lifecycle interpretation, or authenticity evaluation remain explicitly bound to a single synthetic identity and preserve reconstructable lineage sufficient to maintain attribution integrity across reconstruction, transformation, restoration, and cross-context use.

8.2.1 Authority Anchoring Lineage

Requirement

A synthetic identity must maintain an auditable lineage of authority anchoring sufficient to trace authority origin and responsibility across lifecycle transitions.

Interpretation

Authority attachment and transitions must remain traceable such that responsibility can be reconstructed across time and context (NIST, 2023). The origin of authority anchoring begins with the identity's issuing authority as defined by the identity continuity invariants in Section 8.1. Subsequent delegation, transfer, or modification of authority must therefore preserve a continuous chain linking each transition to the original issuing authority or to a clearly attributable successor authority within the governance domain.

Following replication or fork events, authority anchoring lineage must remain explicitly branch-specific; subsequent anchoring transitions on one branch must not be represented as applying retroactively or concurrently to sibling branches.

Each derived branch resulting from replication, cloning, or fork events must therefore retain a distinct identity lineage reference while preserving a verifiable relationship to the originating identity state from which it emerged. Derived identities inherit historical lineage up to the point of divergence but thereafter evolve independently with respect to lifecycle transitions and authority anchoring. Structural continuity therefore applies to each lineage branch individually rather than requiring convergence to a single surviving instance.

Failure Condition

When authority origin cannot be reconstructed, attribution and liability ambiguity emerge.

8.2.2 Non-Silent Lifecycle and Lineage Semantics

Requirement

A synthetic identity must maintain coherent, auditable lifecycle and lineage transitions, including creation, suspension, migration, replication (including cloning and forking), restoration, merge events, and termination. Identity must not silently multiply, resurrect, reset, or migrate without explicit lineage declaration and continuity traceability.

Interpretation

Parallel instantiation is permitted when explicitly represented in lineage semantics. When structural invariants are satisfied, multiple lineage branches derived from a common prior state may each retain continuity validity independent of singular precedence. Forking or cloning does not constitute continuity failure when declared and structurally traceable.

Termination must result in a durable lifecycle tombstone that preserves the historical existence of the identity and prevents identifier reuse. Restoration or reactivation of a terminated identity may occur only as an explicit lifecycle transition that preserves the identity's identifier and historical lineage. Such restoration must remain attributable and auditable and must not obscure the prior termination event.

Material state resets that affect reconstructability of attribution or lineage must be represented as explicit lifecycle events. Silent state erasure, undeclared identity replacement, or implicit resurrection that degrades verifiable continuity constitutes structural failure.

SICF treats identity continuity as a property of lineage coherence rather than singular physical instantiation. Multiple derivative branches may legitimately retain continuity when lifecycle transitions and lineage relationships remain explicitly represented. Structural continuity does not require that only one descendant instance exist at any given time. Determining precedence among branches, resolving authority conflicts, or selecting a canonical successor identity are governance-layer responsibilities rather than properties of identity continuity itself.

Failure Condition

When lifecycle transitions are unrecorded, undeclared, or indistinguishable from replacement, portability failure and vendor-bound identity collapse occur.

8.2.3 Verifiable Continuity and Attribution

Requirement

Identity continuity must be verifiable through durable evidence; the degree of independent verifiability required is determined by the intended assurance level, sufficient to reconstruct lineage and attribute actions across authority chains, mediation layers, and deployment contexts, including evaluated system state and deployed configuration (NIST, 2023).

Evidence Degradation Semantics (Structural)

If continuity evidence is partially lost, corrupted, or rendered non-reconstructable within the declared trust boundary, continuity validity may remain intact only if the remaining evidence is sufficient to reconstruct lifecycle transitions and attribute actions to the identity at the claimed assurance level. Where evidence loss prevents independent reconstruction of lineage or attribution within the boundary, continuity strength collapses to the highest lower level supportable by surviving evidence. If no reconstructable lineage or attribution chain remains, continuity becomes assertion-only and does not satisfy the requirements of this section for external verifiability.

Failure Condition

When continuity cannot be independently reconstructed, auditability failure and attribution instability arise.

The relationship between authority origin, actor attribution, and preserved evidence determines whether governance reconstructability can be maintained. Authority establishes the source of permission, attribution binds actions to an identifiable actor, and evidence preserves the information necessary for later verification. The origin of authority must remain traceable to the issuing authority defined by the identity continuity invariants in Section 8.1, and lifecycle or authority transitions must remain explicitly recorded in order to preserve reconstructable continuity. The interaction of these elements is illustrated in Figure 8-2.

Authority – Attribution – Evidence Relationship

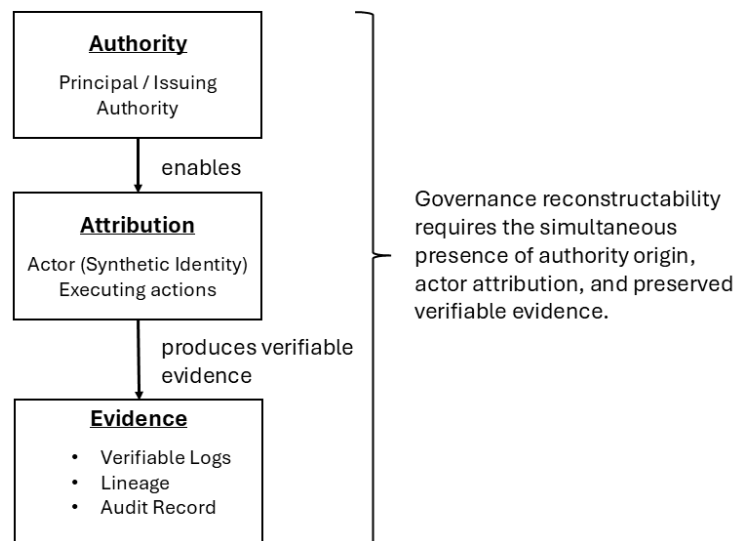


Figure 8-2— Authority–Attribution–Evidence Relationship

Requirement

A synthetic identity must support verifiable authenticity within relevant trust boundaries such that authorized parties can distinguish the legitimate identity from impostors, replayed instances, or unauthorized replicas (NIST, 2020).

Compromise Semantics (Structural)

If the authority anchoring mechanism or authenticity root for a synthetic identity is credibly compromised, continuity must not be treated as high-assurance within the affected trust boundary until a non-silent re-anchoring event is recorded in lineage. Re-anchoring events—including re-keying or replacement of authenticity anchors—must be represented as explicit lifecycle or lineage transitions preserving reconstructable attribution across the pre- and post-compromise boundary. The legitimacy of any such re-anchoring must remain attributable to the issuing authority of the identity or to a governance successor authority operating within the governed namespace defined by the identity continuity invariants in Section 8.1. If compromise cannot be bounded or reconstructably remediated within that boundary, continuity strength collapses to assertion-only within the affected scope. Continuity strength must be evaluated from the earliest credibly inferred compromise point within the declared threat model, and operations occurring between compromise onset and detection must be assessed under the applicable assurance-level degradation semantics rather than presumed to retain pre-compromise strength.

Authenticity verification must incorporate context-bound temporal and lineage-state validation such that previously valid artifacts cannot be replayed outside their original execution context, trust boundary, or declared lifecycle position without detectable discontinuity under the applicable assurance level.

Failure Condition

When unauthorized instances cannot be reliably distinguished from legitimate identity, impersonation and spoofing failure occur.

8.2.4 Persistent Structural Synthetic Classification

Requirement

A synthetic identity must maintain persistent, unambiguous structural classification as synthetic within continuity and attribution semantics, such that responsibility and authority binding cannot be confused with biological or human identity constructs.

Clarification

This requirement constrains structural classification for attribution and authority-binding semantics only and makes no claim regarding moral status, rights, future legal classification, expressive style, personality simulation, or affective behavior.

Structural synthetic classification must remain invariant across embodiment, expressive realism, human identity attestation integration, legal reclassification, or biometric association; no external integration may redefine the identity as biological or human within continuity or attribution semantics.

Recent U.S. judicial reinforcement of the requirement that creative works must originate from human authorship underscores the continued legal distinction between human agency and synthetic output. SICF is architected on the premise that this ontological separation persists for purposes of authority anchoring and attribution semantics. Synthetic identities, regardless of capability, expressive realism, or autonomy, are not treated within this framework as independent human legal subjects, but as structurally synthetic entities anchored to human principal authority.

Failure Condition

When classification becomes ambiguous, responsibility and authority binding destabilize.

Clarification on Derivative Identity Precedence

SICF does not define precedence among legitimate derivative identities. When multiple continuity-preserving branches exist, conflict resolution, authority prioritization, or succession rules are matters for governance-layer frameworks and do not alter structural continuity.

Structural Sufficiency Note

Structural sufficiency within SICF arises from the interaction between the identity continuity invariants defined in Section 8.1 and the continuity requirements defined in this section. The invariants establish the foundational structural conditions necessary for persistent synthetic identity, while the requirements describe the semantic properties that must be preserved in order for continuity to remain reconstructable across time, lifecycle transitions, and deployment contexts. Appendix D identifies structural threat surfaces under which these invariants and requirements may degrade or fail. That appendix does not introduce additional requirements; it stress-tests the invariant and requirement set defined in this section against adversarial, accidental, and systemic discontinuity vectors.

8.2.5 Identity-Bound State and Attribution Constraint

Requirement

State associated with a synthetic identity must be explicitly bound to that identity and remain attributable across all contexts in which it is accessed, maintained, or reconstructed. State that cannot be deterministically associated with a single, continuous synthetic identity does not constitute valid continuity state under SICF.

Derived State Semantics (Structural)

Persistent state alone does not imply continuity. State that is reconstructed within continuity semantics, summarized, transformed, or otherwise derived must retain explicit lineage to its originating identity and prior state to preserve attribution integrity. Derived or transformed state must not be treated as equivalent to original state in contexts requiring verifiable continuity, attribution, or audit reconstruction.

If derived or transformed state results in loss of lineage, ambiguity of attribution, or inability to reconstruct prior state within continuity semantics and the declared trust boundary, continuity validity degrades to the highest level supportable by the remaining reconstructable lineage. Where no reconstructable lineage remains, the resulting state is treated as non-continuity-valid state regardless of apparent persistence or behavioral coherence.

This requirement constrains attribution and lineage semantics only and does not prescribe memory architectures, storage mechanisms, summarization techniques, or implementation strategies.

Failure Condition

When state cannot be deterministically bound to a single continuous synthetic identity or lineage to originating state cannot be reconstructed, attribution integrity fails and continuity becomes non-reconstructable.

9 Terminology and Definitions

The following definitions establish the normative vocabulary of the Synthetic Identity Continuity Framework (SICF). These terms define the structural concepts used throughout the framework and must be interpreted consistently with the continuity and attribution semantics specified herein. Definitions describe conceptual meaning rather than implementation detail.

9.1 Synthetic Identity Continuity Framework (SICF)

The Synthetic Identity Continuity Framework (SICF) is a conceptual schema defining the structural identity invariants and continuity requirements under which a synthetic identity can persist coherently across time, lifecycle transitions, devices, and execution environments.

SICF specifies semantic requirements for continuity, attribution, and authority anchoring lineage independent of model implementation, runtime configuration, hardware embodiment, vendor platform, or physical device.

SICF is governance-neutral and implementation-neutral. It does not prescribe technical architectures, regulatory regimes, certification standards, or ethical doctrine. It defines structural prerequisites upon which downstream governance and application-specific systems may reliably operate.

9.2 Synthetic Actor System (SAS)

The Synthetic Actor System (SAS) is the system-layer architecture that operationalizes identity continuity under governance constraints. It defines how synthetic actors execute actions, maintain and evolve state, and implement lifecycle transitions at runtime across time, context, and infrastructure. SAS enforces identity attribution and governance constraints at execution but does not define identity continuity or authority conditions, which remain the responsibility of SICF and the governance layer, respectively.

9.3 Synthetic Identity

A Synthetic Identity is a persistent semantic continuity construct instantiated through an artificial system that remains structurally coherent across time and transition, independent of any single instance, infrastructure layer, hardware embodiment, device, vendor platform, or model implementation.

A synthetic identity is not equivalent to a deployment instance, runtime process, process identifier, cloud account, executable artifact, or conversational interface. It is the structural continuity referent to which actions, authority relationships, lifecycle transitions, and accountability semantics attach (ISO, 2019; NIST, 2017).

9.4 Identity Identifier

An Identity Identifier is a globally unique identifier bound to a synthetic identity that serves as the primary reference for continuity, lineage, attribution, lifecycle state, and issuer provenance across time, infrastructure environments, and execution contexts.

An identity identifier must remain uniquely associated with a single synthetic identity within the declared trust boundary and must not be reassigned or reused following termination. Identifier continuity enables reconstruction of lifecycle transitions, authority anchoring lineage, and attribution across distributed systems.

9.5 Identity Issuer

An Identity Issuer is an entity authorized within a defined trust boundary to create a synthetic identity and bind a globally unique identity identifier to that identity at the moment of issuance.

Identity issuers establish the initial provenance of identity creation and operate within a governed identifier namespace structure that preserves global uniqueness and traceable issuance authority.

9.6 Identity Continuity

Identity Continuity is the condition under which a synthetic identity remains the same identity across lifecycle transitions, migration, replication, restoration, device reassignment, or environmental change.

Continuity does not require immutability of internal state, model weights, memory contents, embodiment, or deployment environment. It requires preservation of identity invariants, lineage, authority anchoring, and verifiable linkage across time such that past and present states can be coherently attributed to the same identity construct.

Replacement or modification of model architecture, training weights, inference stack, or implementation mechanism does not by itself redefine identity continuity, provided lifecycle and authority anchoring invariants remain satisfied.

9.7 Persistence

Persistence refers to the temporal extension of a synthetic identity beyond bounded tasks or session-scoped interaction.

Persistence describes duration, while identity continuity describes the structural coherence of that persistence across change.

A system may persist temporally without preserving continuity; continuity specifies the conditions under which persistence remains stable.

Persistence may include continuation across:

- Devices
- Vendors
- Runtime environments
- Execution architectures
- Hardware embodiments
- Jurisdictions

9.8 Instance

An Instance is a specific runtime instantiation of a synthetic identity within a defined execution environment.

Multiple instances may exist in parallel; when instances evolve independently from a common prior state, the divergence constitutes a fork and must be explicitly represented within lineage semantics. An instance is not equivalent to a synthetic identity, and identity continuity may persist across changes in instance, environment, hardware, or infrastructure.

9.9 Agency

Agency describes what a system can do and refers to functional capability to perceive, decide, and act within defined constraints.

Agency does not imply legal status, moral standing, or personhood.

Synthetic identity describes which continuity-bearing construct those actions attach to across time. An agent operates through a synthetic identity, so a synthetic identity may persist even as agency capabilities evolve, expand, or contract.

Agency is behavioral capacity. Synthetic identity is structural continuity. The two are related but not equivalent.

9.10 Principal

A Principal is a human or institutional entity from which delegated authority for a synthetic identity originates.

Institutional entities may include corporations, government bodies, regulated agencies, nonprofit organizations, courts, or other legally constituted organizations.

A principal may delegate authority to a synthetic identity. Authority relationships may evolve over time and may include multiple principals or institutional anchors. This definition does not imply ownership, legal personhood, or exclusive control.

9.11 Authority Anchoring

Authority Anchoring refers to the structural binding between a synthetic identity and the principal(s) from which its delegated authority originates. Authority anchoring establishes the origin and lineage of responsibility associated with actions performed by the identity.

Authority anchoring defines the source and traceable lineage of responsibility. It does not define the protocols by which authority is delegated, scoped, enforced, or revoked; those mechanics belong to governance-layer frameworks. Authority anchoring is distinct from operational control, infrastructure ownership, or technical access.

Authority anchoring lineage must remain reconstructable across lifecycle transitions if identity continuity is to remain coherent. Transitions in authority anchoring do not require identity termination; a synthetic

identity may persist across changes in principal provided such transitions are explicitly represented within lineage semantics.

Authority anchoring specifies the traceable origin of responsibility only. Delegation rules, scope constraints, enforcement mechanisms, precedence, and revocation mechanics remain governance-layer concerns. Conflicting, overlapping, or revoked principal relationships do not invalidate identity continuity provided anchoring lineage remains reconstructable (NIST, 2023).

9.12 Attribution

Attribution refers to the structured linkage between a synthetic identity and actions across authority chains and execution environments.

Within SICF, attribution refers specifically to action-binding semantics — the ability to determine which synthetic identity performed or authorized a given action and under what authority conditions.

Attribution requires structural linkage between identity, authority anchoring, and evidentiary continuity (NIST, 2023).

9.13 Lifecycle

Lifecycle refers to the structured sequence of identity-relevant states and transitions a synthetic identity may undergo across its existence.

Lifecycle events include, but are not limited to:

- Creation
- Suspension
- Migration
- Replication
- Restoration
- Fork or merge events
- Device reassignment
- Termination

Lifecycle transitions must not occur silently if continuity is to remain coherent.

A silent lifecycle transition refers to any state-altering event that changes the operational substrate, authority anchoring context, lineage position, identifier binding, or active instance topology of a synthetic identity without explicit declaration at the identity layer. Such transitions may include undeclared restoration from backup, unrecorded replication, implicit forks from desynchronized distributed execution, identifier reassignment, rollback without lineage representation, or substrate substitution without a migration record.

Replication refers to any event in which a synthetic identity produces one or more derivative instances derived from a prior state and may take multiple forms:

- Cloning — state-identical duplication at the moment of replication
- Forking — lineage divergence from a common prior state

Cloning is a subset of replication. Forking refers to divergence within lineage semantics.

Forking includes both intentional replication events and divergence arising from distributed or desynchronized operation. When multiple instances of a synthetic identity evolve independently from a common prior state, the divergence constitutes a fork and must be explicitly represented within lineage semantics.

Both cloning and forking must be explicitly represented within identity lineage.

Restoration refers to re-instantiation from a prior recorded state of a synthetic identity. Restoration occurring while an identity remains active constitutes a fork and must be explicitly represented within lineage semantics.

Migration across substrates must preserve authority anchoring lineage such that continuity claims remain reconstructable independent of the operational state of the originating substrate. Subsequent compromise, decommissioning, or adversarial control of a prior substrate must not retroactively invalidate continuity of a successfully recorded migration event within the declared trust boundary. Competing lineage claims originating from a compromised or superseded substrate must be treated as fork events subject to declared assurance-level detection and integrity validation requirements.

Merge refers to a lifecycle event in which two or more lineage branches derived from a common synthetic identity are recombined into a single continuity-bearing identity. A merge event must preserve reconstructable lineage of each contributing branch and must not erase or obscure prior divergence. A merge must not overwrite, collapse, or reinterpret authority anchoring or attribution history of contributing branches; divergent histories must remain independently reconstructable after the merge.

A merge does not imply retroactive identity singularity. It represents explicit convergence following divergence within recorded lineage semantics.

Termination refers to a lifecycle event in which a synthetic identity's authority anchoring and agency cease and no further lifecycle transitions are permitted. Termination closes active continuity while preserving reconstructable lineage and attribution.

A terminated identity remains historically attributable within defined trust and assurance boundaries. Termination is irreversible for the terminated identity; any re-instantiation derived from prior state following termination constitutes a new synthetic identity with explicit lineage reference to the predecessor.

A terminated identity identifier must not be silently reactivated, reassigned, or resumed; any post-termination operational instance must be explicitly instantiated as a distinct synthetic identity within lineage semantics.

Identity identifiers must be uniquely bound within the declared trust boundary such that reuse, reassignment, or namespace collision cannot simulate continuity or authority anchoring without explicit lineage reference to the predecessor identity. Identifier uniqueness is a structural requirement for

continuity integrity; namespace ambiguity or cross-boundary identifier reuse must be interpreted as a fork or new instantiation event requiring explicit lineage representation.

9.14 Lineage

Lineage refers to the reconstructable historical chain of lifecycle transitions, authority changes, and structural state evolution associated with a synthetic identity.

Lineage records how an identity has transitioned across time, including migration, replication, restoration, merge, delegation, device reassignment, or termination events.

Lineage is distinct from internal memory.

9.15 Continuity Requirements

Continuity Requirements are the minimal structural conditions that must hold for identity continuity to remain coherent under SICF.

These requirements are defined normatively in Section 8.2 and are independent of implementation mechanism or governance model.

9.16 Continuity Validity

Continuity Validity refers to the condition under which the structural invariants defined in Section 8.1 remain satisfied and the continuity requirements defined in Section 8.2 remain structurally preservable. When continuity validity holds, identity remains structurally coherent independent of evidentiary strength or governance context.

9.17 Continuity Strength

Continuity Strength refers to the degree to which the identity continuity invariants defined in Section 8.1 and the continuity requirements defined in Section 8.2 are independently verifiable within defined trust boundaries and declared threat models.

Continuity strength may be expressed through graded assurance levels and does not alter structural continuity validity.

9.18 Trust Boundary

A Trust Boundary is the defined scope within which authenticity, continuity, authority, and attribution claims are evaluated and verified.

Trust boundaries may correspond to institutional, organizational, technical, jurisdictional, contractual, or other formally recognized domains. Within a given trust boundary, authorized parties must be able to distinguish legitimate synthetic identities from unauthorized, spoofed, replayed, or improperly replicated instances.

Trust boundaries define the contextual perimeter in which continuity and authenticity are assessed. They do not prescribe specific verification mechanisms or technical controls. Different trust boundaries may

apply different assurance levels, provided that continuity and attribution remain reconstructable within the relevant domain (NIST, 2020).

10 Layered Model of the SICF Ecosystem

The Synthetic Identity Continuity Framework (SICF) functions as the foundational layer in a multi-layered ecosystem of synthetic systems. Its purpose is not to govern behavior, enforce policy, or prescribe implementation architecture. Its function is to define the structural conditions under which a synthetic identity persists coherently across time and transition.

To preserve conceptual clarity and prevent scope conflation, SICF is positioned as Layer 0 within a layered model as illustrated in Figure 10-1.

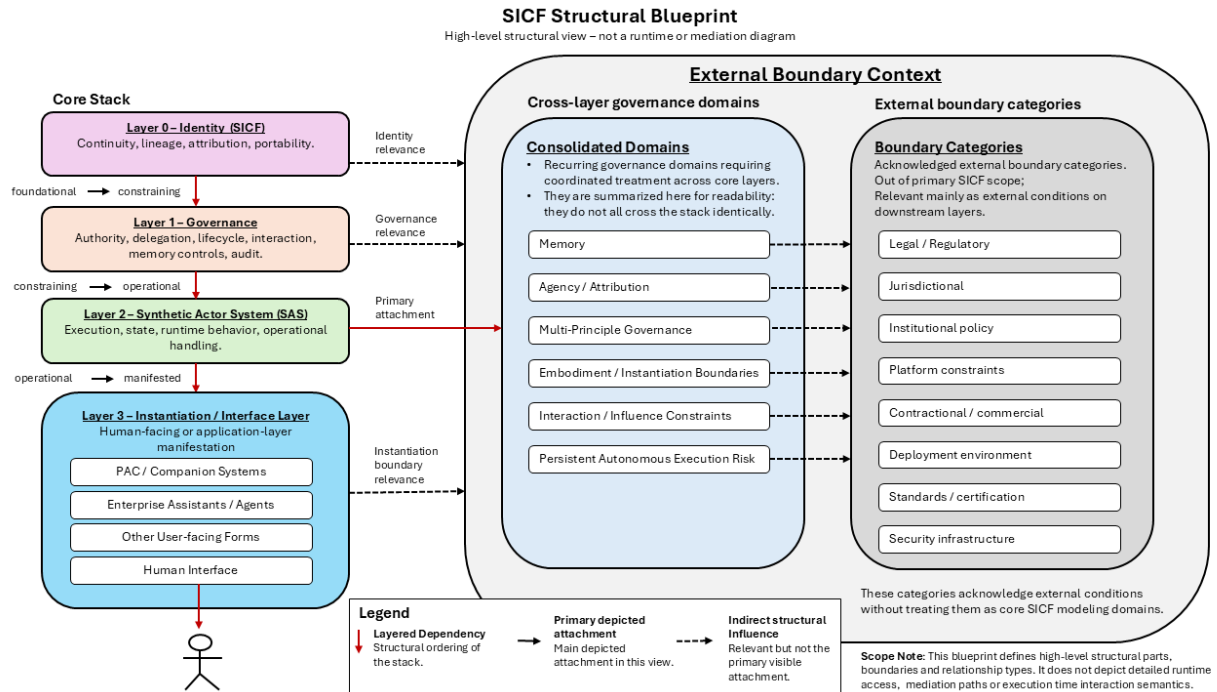


Figure 10-1 – SICF Structural Blueprint (High-Level Layered Ecosystem View)

10.1 Layer 0 — Identity Continuity Layer (SICF)

Layer 0 defines the minimal structural invariants required for persistent synthetic identity. These invariants establish the foundational identity conditions under which continuity can exist across time, system transitions, and deployment environments. The invariants defined in Section 8.1 include global identifier uniqueness, identifier non-reuse, verifiable issuer provenance, governed namespace allocation, durable lifecycle state, state coherence across distributed systems, and explicit lifecycle and authority transitions.

These invariants together ensure that synthetic identities remain uniquely identifiable, historically traceable, and structurally coherent across migrations, forks, restorations, and governance transitions. The continuity requirements defined in Section 8.2 build upon these invariants to define the semantic conditions under which lineage, attribution, authority relationships, authenticity within declared trust

boundaries, persistent structural synthetic classification, and identity-bound state and attribution semantics remain reconstructable across identity-relevant change.

Layer 0 does not define governance policy, behavioral constraints, risk management rules, or operational controls. It defines only the structural conditions under which identity continuity exists.

Layer 0 is intentionally minimal. Future governance architecture (Layer 1), Synthetic Actor Systems (Layer 2), and application instantiations (Layer 3) may build upon these invariants and continuity requirements, but such downstream layers must not redefine, narrow, weaken, or substitute for the normative identity continuity conditions established in Sections 8.1 and 8.2. Downstream elaboration may constrain behavior, authority, execution, state handling, or deployment context, but it must preserve the continuity semantics defined at Layer 0 as the upstream structural precondition for all later layers.

Without Layer 0, higher-order governance mechanisms operate on unstable identity substrates.

Layer 0 invariants are substrate-agnostic by design. They must survive model substitution, runtime migration, hardware replacement, embodiment transitions, jurisdictional redeployment, and cloud-to-edge re-instantiation without semantic redefinition. Continuity anchored to a specific model weight configuration, infrastructure provider, hardware root, or embodiment artifact is not structural identity continuity but implementation binding. Substrate dependence introduces implicit reset conditions that undermine lineage integrity, attribution stability, and long-term accountability. For this reason, identity continuity within SICF must remain logically independent of any specific runtime implementation even as cognitive, expressive, or physical implementations evolve.

10.2 Layer 1 — Governance Architecture

Layer 1 defines governance architecture for persistent synthetic identities. Governance architecture specifies the structured control plane within which authority is exercised, delegated, time-bound, scope-constrained, monitored, transitioned, and revoked. Delegation may be bounded by duration, scope, capability class, assurance requirement, or contextual trigger. Expiration, modification, or revocation of delegated authority does not redefine the underlying synthetic identity nor alter its continuity invariants; such changes must be explicitly represented within authority anchoring lineage semantics. Governance architecture therefore governs how authority flows and terminates over time without mutating identity continuity or retroactively altering attribution history.

Layer 1 does not redefine identity continuity. It operates upon the invariants defined in Layer 0 and presumes their stability. Governance architecture determines how authority flows across principals, how authority conditions are applied to lifecycle transitions, how assurance levels are required and validated, and how disputes or ambiguity are resolved within defined trust contexts.

Governance architecture is domain-agnostic. It does not prescribe enterprise, public-sector, companion, or mission-critical deployment characteristics. Instead, it establishes the structural mechanisms through which those application instantiations may configure authority, risk tolerance, verification rigor, and operational constraints.

Layer 1 therefore functions as the structural bridge between identity continuity and system operation. Synthetic Actor Systems (Layer 2) implement and enforce governance architecture at runtime, while application instantiations (Layer 3) configure and apply governance within specific operational contexts without redefining its structural logic.

Governance defines the conditions under which actions are permitted, including authority validation, delegation constraints, and audit requirements. It does not define how actions are executed, how state is maintained, or how lifecycle transitions occur. These responsibilities are implemented within the Synthetic Actor System (SAS) layer, which operationalizes governance constraints at runtime while preserving identity continuity as defined by SICF.

10.3 Layer 2 — Synthetic Actor System (SAS)

The Synthetic Actor System (SAS) defines the system-layer architecture responsible for operationalizing identity continuity under governance constraints. It provides the runtime environment in which persistent synthetic actors exist and operate across time, context, and infrastructure.

SAS defines how synthetic actors execute actions, maintain and evolve state, and implement lifecycle transitions at runtime. It establishes execution models supporting deferred and asynchronous operation, state persistence across sessions and environments, and lifecycle consistency across instantiation, suspension, and termination states.

Within this layer, identity continuity and governance constraints are enforced at execution. SAS ensures that actions are performed under valid authority conditions at the time of execution, that attribution remains consistent, and that state evolution remains coherent with both continuity requirements and governance constraints across distributed and multi-context environments.

SAS does not define identity continuity or authority conditions. Identity continuity remains the responsibility of SICF, and governance defines permissible actions and authority constraints. SAS implements these definitions at runtime, providing the system-level mechanisms through which persistent, attributable, and auditable synthetic actors operate. Application instantiations are constructed on top of SAS and inherit its execution, state, and lifecycle properties.

10.4 Layer 3 — Application Instantiations

Layer-3 Application Instantiations represent classes of persistent synthetic identities instantiated within Synthetic Actor Systems (Layer 2) that inherit Layer-0 continuity invariants and are governed by Layer-1 Governance Architecture. Layer 3 does not redefine continuity semantics; it applies them within concrete deployment contexts.

Application instantiations operationalize synthetic identity across diverse domains while remaining structurally bound to the identity continuity invariants and requirements defined in Sections 8.1 and 8.2. These systems may operate across heterogeneous models, distributed infrastructures, and evolving configurations, yet continuity validity remains anchored to authority lineage, non-silent lifecycle semantics, verifiable attribution, authenticity within declared trust boundaries, persistent structural synthetic classification, and identity-bound state and attribution semantics. Application diversity, domain specialization, and implementation change do not alter the requirement that continuity-relevant state

remain explicitly bound to a single synthetic identity with reconstructable lineage sufficient to preserve attribution integrity across reconstruction, transformation, and cross-context use.

Layer-3 instantiations may include enterprise automation agents, public-sector service agents, mission-critical operational agents, persistent AI companions, and multi-principal synthetic systems. These categories are illustrative rather than exhaustive and do not constitute normative prescriptions.

Layer 3 concerns how synthetic identities are instantiated and governed in real-world operational contexts. Application diversity does not imply identity variability; continuity semantics remain invariant across instantiation classes.

10.5 Implementation Layer (Runtime and Infrastructure)

Implementation environments represent the technical realization of the Synthetic Actor System (SAS) layer and include:

- Model architectures and weights
- Runtime orchestration platforms
- Containerized or distributed execution environments
- Edge or embedded hardware
- Cloud and vendor infrastructure
- Interface surfaces and embodiment layers

Implementation mechanisms may change without altering synthetic identity continuity, provided Layer-0 structural invariants are preserved and Layer-1 governance architecture remains coherently applied within Synthetic Actor Systems (Layer 2). Execution of cognitive processes may occur partially or entirely across distributed environments, including cloud infrastructure, edge nodes, embedded systems, or sovereign compute domains; identity continuity must not depend upon a singular execution locus or centralized processing environment.

SICF is independent of specific technical implementations and remains valid across architectural evolution. Implementation environments serve as the underlying infrastructure through which Synthetic Actor Systems are realized, but do not constitute a separate architectural layer within the SICF model.

10.6 Legal and Regulatory Context (Orthogonal Layer)

Legal and regulatory regimes operate in relation to Layer-1 governance architecture, Synthetic Actor Systems (Layer 2), and application instantiations (Layer 3) but are not defined by SICF. Law may interpret synthetic identity constructs, delegation structures, and liability chains, but legal classification does not alter structural continuity semantics.

SICF defines structural identity prerequisites that may support legal interpretation without presupposing it.

In this framework, orthogonal denotes a domain that intersects with, constrains, or influences the layered architecture without being hierarchically subordinate to it or redefining its structural invariants.

10.7 Layer Separation and Invariance Principles

Each layer serves a distinct function:

- Layer 0 defines structural identity continuity invariants.
- Layer 1 defines Governance Architecture, including authority control, delegation, assurance mechanisms, and policy enforcement structures.
- Layer 2 defines the Synthetic Actor System (SAS), which provides the runtime environment in which persistent synthetic identities execute, maintain state, and operate under Layer-1 governance constraints.
- Layer 3 defines application instantiations, which configure and apply governance within domain-specific operational contexts built on Synthetic Actor Systems.
- Implementation environments realize the Synthetic Actor System through runtime infrastructure, execution platforms, and deployment substrates.
- Legal and regulatory regimes interpret, constrain, and regulate deployment contexts without redefining structural continuity invariants.

No layer substitutes for another. Governance does not redefine structural continuity invariants; application diversity does not weaken invariants; implementation variability does not determine continuity validity; and legal interpretation does not retroactively alter structural continuity semantics.

Layer separation is a structural requirement for preserving identity continuity across technological evolution, jurisdictional variation, and deployment diversity.

Layer Invariance Principle

Layer 0 structural continuity semantics—comprising the identity continuity invariants defined in Section 8.1 and the continuity requirements defined in Section 8.2—must remain satisfied under any change in governance model, implementation architecture, embodiment, expressive behavior, personality traits, legal classification, or external identity integration within the scope and assumptions defined in Section 7 and within declared trust boundaries. Downstream layers may operationalize, constrain, or extend behavior around these semantics, but they must not reinterpret, narrow, or substitute for the normative continuity conditions established in Sections 8.1 and 8.2.

Personality traits, expressive realism, embodiment characteristics, and inference latency are not components of Identity Continuity. They may vary across Governance Architecture (Layer 1), Synthetic Actor Systems (Layer 2), implementation environments, and application instantiations without affecting Layer 0 structural invariants.

Radical behavioral change, personality drift, value evolution, or expressive mutation does not invalidate identity continuity so long as the Layer 0 structural invariants defined in Section 8.1 remain satisfied and the continuity requirements defined in Section 8.2 remain preserved. Changes in behavior, expression, or modeled disposition may be operationally significant within governance, application, or implementation contexts, but they do not by themselves alter continuity validity under SICF.

Interoperation with external human identity attestation or biometric verification systems must preserve structural separation between synthetic and human identity constructs. Layer 0 continuity invariants must not depend upon a single human identity authority, centralized registry, or jurisdiction-specific attestation provider. Continuity must remain valid under multiple providers, jurisdictional variation, revocation events, and consent controls defined within governance layers.

Upward Dependency Principle

Dependency flows upward. Governance architecture (Layer 1), Synthetic Actor Systems (Layer 2), application instantiations (Layer 3), implementation mechanisms, and legal interpretations depend upon identity continuity (Layer 0). Synthetic Actor Systems depend upon governance architecture, and application instantiations depend upon Synthetic Actor Systems and governance architecture. Identity continuity does not depend upon governance design, system implementation, application configuration, or legal classification.

Delegation mechanics, authority transfer protocols, trust calibration models, and recovery logic are governance-layer concerns. Layer 0 provides the persistent identity substrate required for delegation to be meaningful but does not define how authority is transferred, constrained, or revoked.

Cross-layer conflation introduces structural instability. Maintaining disciplined layer separation preserves conceptual clarity and prevents redefinition of identity continuity by downstream governance, implementation, or legal interpretation.

11 Related Work and Conceptual Positioning

SICF does not emerge in isolation. Questions of identity, attribution, continuity, and authority have been studied across multiple domains, including digital identity systems, distributed systems theory, software supply chain security, privacy regulation, and governance frameworks for artificial intelligence (European Parliament and Council, 2024; NIST, 2023; OECD, 2019; W3C, 2022). However, these domains address adjacent but structurally distinct problems.

SICF defines a continuity layer not fully specified within existing frameworks.

11.1 Digital Identity and Authentication Systems

Digital identity systems—including public key infrastructures (PKI), federated identity models, decentralized identifiers (DIDs), and self-sovereign identity (SSI) architectures—rely on cryptographic authentication, credential verification, and trust establishment mechanisms (Diffie & Hellman, 1976; W3C, 2022).

Persistent identifier infrastructures provide an additional example of durable identity constructs in large-scale digital ecosystems. Systems such as the Digital Object Identifier (DOI) framework and the Handle System demonstrate that globally unique identifiers can remain stable across infrastructure change, repository migration, and long-lived reference chains. Similarly, researcher identity systems such as ORCID maintain persistent identity anchors across institutional and organizational transitions. These systems illustrate the feasibility and value of durable identifier persistence. However, they address documents, datasets, or human identity records rather than executable synthetic actors operating across runtime environments. SICF extends the concept of durable identifiers to persistent synthetic identities whose lifecycle includes fork, restoration, migration, and parallel instantiation events.

Operational identity infrastructures also depend on the lifecycle management of trust anchors and credential issuance chains. In practice, identity systems must address certificate issuance, revocation, expiration, and trust-root governance across long-lived operational environments. Discussions within identity infrastructure communities emphasize that maintaining trust in identity systems requires disciplined lifecycle management of these anchors and verification chains over time. While such mechanisms preserve trust in identity assertions, they remain focused on credential validity rather than the structural continuity of the identity entity itself.

These mechanisms establish who or what is presenting credentials at a given moment but do not define lifecycle continuity semantics for persistent synthetic identities across fork, merge, restoration, replication, or termination events. Authentication supports authenticity within trust boundaries; it does not define cross-instance identity continuity.

Emerging digital infrastructure initiatives increasingly emphasize persistent identifiers as foundational elements of large-scale information ecosystems. Research and data infrastructures, for example, rely on persistent identifier systems to maintain stable attribution and referenceability across distributed repositories and long-lived digital artifacts. These efforts reinforce the importance of durable identifier semantics in complex digital environments. However, while persistent identifiers preserve referential

stability, they do not address the broader lifecycle continuity challenges posed by synthetic agents operating across replication, migration, and infrastructure transitions. SICF extends beyond identifier persistence by defining structural continuity invariants capable of preserving identity semantics across such transformations, ensuring that identity lineage and authority binding remain verifiable despite replication, migration, or restoration events.

Synthetic identity under SICF is distinct from master data management (MDM) constructs. While both require canonical identifiers and cross-system consistency, master data represents passive referential entities. Synthetic identity, by contrast, incorporates lifecycle semantics, authority anchoring, and continuity invariants governing identity-relevant transitions such as replication, migration, restoration, and termination across persistent instantiations.

11.2 Distributed Systems and Consistency Models

Distributed systems research addresses replication, divergence, consensus, and convergence under partition (Lamport, 1978).

While these models explain how state evolves across distributed environments, they do not define normative continuity semantics for synthetic identity across fork and merge events. SICF adopts distributed-systems vocabulary while defining structural identity invariants independent of specific consistency mechanisms.

11.3 Software Supply Chain and Provenance Frameworks

Provenance and attestation systems focus on artifact integrity, version traceability, and build reproducibility (Open Source Security Foundation, 2023; Torres-Arias, Afzali, Dolan-Gavitt, & Cappos, 2019). These frameworks attach evidentiary history to software artifacts rather than to persistent synthetic identities operating across lifecycle transitions. SICF extends continuity semantics from artifact provenance to identity-level lineage, authority anchoring, and state evolution.

Modern software supply-chain security initiatives increasingly emphasize artifact provenance and lifecycle transparency as foundational infrastructure properties. Frameworks such as Supply-chain Levels for Software Artifacts (SLSA) formalize build provenance, dependency lineage, and verifiable artifact history across complex software pipelines (Open Source Security Foundation, 2023). Similarly, Software Bill of Materials (SBOM) standards provide structured disclosure of software component composition and dependency relationships, enabling traceability and risk assessment across distributed software ecosystems (National Telecommunications and Information Administration, 2021). These initiatives reflect a broader architectural shift toward explicit provenance and lineage tracking in complex distributed systems. SICF extends this structural principle to persistent synthetic identity, requiring continuity semantics capable of preserving identity lineage and attribution across migration, replication, and long-lived operational environments.

Transparency and tamper-evident logging infrastructures further reinforce the architectural shift toward verifiable lineage in distributed systems. Certificate Transparency and similar append-only logging mechanisms demonstrate how cryptographic evidence chains can expose unauthorized changes, mis-issuance events, or integrity violations across large-scale identity and trust infrastructures (Laurie,

Langley, & Kasper, 2013). These systems illustrate how durable evidence records can support independent verification of historical state transitions. SICF applies a similar evidentiary principle at the identity layer, requiring lifecycle transitions affecting synthetic identity continuity—such as replication, migration, restoration, or termination—to remain externally verifiable through durable lineage records.

11.4 AI Governance Frameworks

AI governance frameworks define policy controls, oversight mechanisms, compliance obligations, and accountability processes (NIST, 2023; OECD, 2019; European Parliament and Council, 2024). Such frameworks presuppose stable identity constructs in order to assign responsibility and evaluate behavior. SICF defines the structural continuity layer upon which governance operates.

Model-level alignment training, safety tuning, or internally defined behavioral “constitutions” shape outputs but do not establish cross-instance identity continuity, lineage semantics, or portability across migration and re-instantiation. These operate at behavioral or governance layers rather than at the identity layer.

11.5 Distinguishing Contribution of SICF

SICF introduces a minimal structural model that:

- Defines synthetic identity as a continuity-bearing construct independent of implementation
- Formalizes lifecycle semantics including fork, merge, restoration, and termination
- Separates identity continuity from governance authority
- Establishes invariance across migration, embodiment, and legal classification
- Identifies structural failure modes arising from undefined continuity semantics
- Positions synthetic identity continuity as a structural infrastructure layer that integrates concepts from persistent identifier systems, distributed systems lineage, and software provenance frameworks while extending them to persistent synthetic actors operating across heterogeneous runtimes

SICF specifies semantic conditions required before governance, compliance, or implementation architectures can operate coherently.

11.6 Privacy and Personally Identifiable Information (PII) Frameworks

Privacy frameworks and PII regulatory regimes govern the collection, storage, processing, and protection of information relating to identifiable individuals (European Parliament and Council, 2016; NIST, 2020). These frameworks define data classification standards, consent requirements, access controls, and subject rights.

SICF does not regulate or classify data. It defines structural continuity semantics for synthetic identities independent of whether associated operational data constitutes personally identifiable information.

Persistent synthetic identities may accumulate, retain, or act upon PII depending on deployment context. Data protection obligations apply at the governance and implementation layers. SICF neither prescribes

nor replaces privacy regulation; it provides identity continuity invariants upon which privacy-compliant governance models may operate.

Continuity metadata itself may be operationally sensitive depending on deployment context; confidentiality, access control, and disclosure policies are governed by Layer-1 mechanisms and applicable regulatory regimes rather than by Layer-0 continuity invariants.

11.7 Zero Trust and Identity-Centric Security Models

Zero Trust architectures and identity-centric security models emphasize continuous verification, context-aware authorization, and minimization of implicit trust within networks and systems (NIST, 2020; NIST, 2023).

While Zero Trust frameworks depend upon reliable identity signals, they focus on access control decisions and security posture. They do not define structural continuity semantics for synthetic identities across lifecycle transitions, replication, restoration, fork, merge, or termination events.

SICF defines identity continuity invariants that Zero Trust and identity-centric security architectures may rely upon but do not themselves specify.

11.8 Hardware Roots of Trust and Cryptographic Anchoring

Modern security architectures often rely upon hardware-backed keys, trusted platform modules (TPMs), secure enclaves, and cryptographic attestation mechanisms to establish authenticity and integrity (TCG, 2019).

Such mechanisms may support aspects of SICF requirements—particularly authenticity within trust boundaries and evidentiary continuity. However, hardware anchoring and cryptographic identity proofs are implementation mechanisms. They do not, by themselves, define identity continuity semantics across distributed instantiation, lifecycle transitions, or authority evolution.

SICF remains independent of specific hardware or cryptographic mechanisms.

11.9 Agentic and Multi-Agent Architectures

Contemporary AI systems increasingly employ agent-based architectures, orchestration frameworks, tool-using agents, and federated multi-agent coordination.

These systems define behavioral capabilities, task delegation structures, and coordination protocols among agents. They do not define structural continuity semantics for persistent synthetic identity across fork, merge, restoration, or termination events.

SICF operates beneath agent architectures. An agent may operate through a synthetic identity as defined by SICF, but agent capability expansion or orchestration complexity does not redefine identity continuity.

11.10 Model Governance and MLOps Traceability

Model governance frameworks, including model registries, training-data lineage systems, version control practices, and MLOps pipelines, focus on artifact traceability and model lifecycle management.

These systems track model versions, training datasets, deployment configurations, and performance metrics. They attach provenance to artifacts rather than to persistent synthetic identities operating across distributed runtime environments.

SICF distinguishes between artifact provenance and identity continuity. Model lineage and identity lineage are related but structurally distinct constructs.

11.11 Ledger-Based Identity and Distributed Anchoring

Some identity frameworks employ distributed ledger technologies or immutable logging systems to anchor identity events and lifecycle transitions (Nakamoto, 2008).

Such mechanisms may support continuity verification or lineage recording. However, ledger infrastructure is an implementation choice rather than a defining characteristic of synthetic identity continuity.

SICF does not require distributed ledger infrastructure. It defines semantic continuity conditions independent of specific anchoring technologies.

12 Assurance Levels and Identity Strength Taxonomy

Identity continuity under SICF is not merely a binary property. Continuity may be preserved at varying levels of structural assurance depending on evidentiary durability, verification scope, and trust boundary rigor.

Assurance Levels defined in this section are independent of the architectural Layers defined in Section 10. Layers describe structural separation of concerns within the SICF ecosystem; assurance levels describe evidentiary strength of continuity verification within a given trust boundary.

This section defines a graded taxonomy of continuity assurance without prescribing implementation mechanisms.

12.1 Rationale for Assurance Gradation

Sections 8.1 and 8.2 define the structural foundations necessary for identity continuity. Section 8.1 specifies the invariants that must always hold for persistent synthetic identity, while Section 8.2 defines the continuity requirements necessary to preserve reconstructable lineage, authority binding, attribution, lifecycle semantics, authenticity within declared trust boundaries, persistent structural synthetic classification, and identity-bound state and attribution semantics. This section introduces a taxonomy describing the evidentiary strength with which those invariants and requirements can be verified. Systems may satisfy structural continuity while differing significantly in the degree of independent verifiability supporting continuity claims.

Graded assurance models are common in security and identity standards, which distinguish between structural validity and the strength of verification evidence under defined trust contexts (NIST, 2017; NIST, 2023).

A synthetic identity operating within a single enterprise boundary may require different assurance characteristics than a synthetic identity operating across jurisdictions, vendors, or mission-critical infrastructure. Assurance levels therefore describe evidentiary strength within a defined trust boundary rather than intrinsic properties of the identity itself.

SICF distinguishes between continuity validity and continuity strength. Continuity validity refers to whether the structural invariants defined in Section 8.1 remain satisfied and the continuity requirements defined in Section 8.2 remain structurally preservable. Continuity strength refers to how robustly and independently those invariants and requirements can be verified within a defined trust boundary.

Some frontier AI governance frameworks increasingly employ internal capability-tier classifications that activate controls at defined performance or risk thresholds. Such tiered models implicitly presume persistent system identity, version lineage, and enforceable attribution across upgrades and deployment contexts. SICF provides the structural identity layer necessary for such threshold-based governance to remain coherent across migration, fork, and jurisdictional transition.

Escalation across capability thresholds must preserve identity continuity semantics, such that governance activation does not imply identity redefinition. Identity continuity must also support version-specific incident reconstruction, including fork traceability and immutable audit linkage.

This distinction allows structural identity continuity to remain minimal and technology-neutral while still supporting high-assurance deployments. Systems operating within bounded or low-risk environments may satisfy continuity requirements with limited independent verification, while mission-critical deployments may require stronger evidence preservation, tamper detection, and cross-boundary verification capabilities. Assurance gradation therefore separates the structural existence of continuity from the evidentiary strength with which continuity claims can be independently validated.

12.2 Structural Assurance Levels

The following levels define increasing assurance strength. These levels are descriptive and may be incorporated within Layer-1 governance architecture and applied by Layer-3 application instantiations through the Synthetic Actor System (Layer 2).

Level 0 — Implicit Continuity (Non-Conformant)

Continuity is assumed but not formally represented.

- Lifecycle transitions may occur without recorded lineage.
- Authority anchoring may be informal or undocumented.
- Identity may be bound to infrastructure or vendor context.

This level does not satisfy the identity continuity invariants defined in Section 8.1 or the continuity requirements derived from those invariants in Section 8.2.

Level 1 — Internal Continuity Assertion

Continuity invariants are defined and internally recorded but not independently verifiable outside the operating environment.

- Lifecycle events are logged.
- Authority anchoring transitions are represented.
- Authenticity mechanisms exist within a bounded environment.

Verification is limited to internal system control.

Level 2 — Bounded External Verifiability

Continuity invariants are verifiable within a defined trust boundary by authorized third parties.

- Lifecycle lineage is reconstructable.
- Authority anchoring transitions are traceable.
- Authenticity claims are testable within the boundary.

This level supports enterprise or institutional audit contexts.

Level 3 — Cross-Boundary Portability and Verification

Continuity invariants remain verifiable across vendor transitions, infrastructure migration, or jurisdictional change.

- Identity lineage survives platform change.
- Authority anchoring history remains reconstructable after migration.
- Authenticity claims are not infrastructure-bound.

This level is intended to support continuity across systemic transition.

Level 4 — High-Assurance Continuity (Declared Threat Model)

Continuity invariants are resilient against adversarial manipulation within explicitly declared threat models.

Assurance claims tied to explicit threat models align with established security engineering practice, in which system guarantees are evaluated relative to declared adversarial capabilities and attack surfaces rather than absolute promises (Shostack, 2014; NIST, 2018).

A system claiming Level 4 assurance must:

- Declare the threat model under which continuity claims are evaluated.
- Define the adversarial capabilities considered (e.g., insider compromise, replay, fork manipulation, credential theft, infrastructure tampering).
- Demonstrate that lifecycle events, lineage integrity, authority anchoring, and authenticity boundaries remain detectably resistant within that threat model.
- Enable independent reconstruction and tamper-evident detection consistent with declared assumptions.
- Demonstrate that fork and merge events cannot be used to obscure, overwrite, or launder compromised authority anchoring or attribution history within the declared threat model.
- Demonstrate that lineage recording, lifecycle event logging, and authority anchoring state transitions are resistant to unauthorized suppression, reordering, or silent modification by actors operating within the declared insider threat model.

High assurance is not absolute; it is defined relative to a transparent and explicitly bounded threat model. Without declared threat assumptions, Level 4 claims are undefined. This level is appropriate for mission-critical, public-sector, financial, safety-critical, or high-liability deployments.

12.3 Assurance Scope and Governance Separation

Assurance levels describe evidentiary strength and verifiability characteristics. They do not prescribe specific cryptographic schemes, hardware anchors, logging standards, or compliance frameworks.

This separation mirrors established distinctions between architectural security properties and governance-layer compliance regimes (NIST, 2023; ISO, 2022).

Layer-1 governance architecture may define:

- Required assurance levels for specific deployment classes
- Sector-specific threat model expectations
- Conformance verification procedures
- Audit and certification methodologies

SICF remains independent of those prescriptions.

12.4 Continuity Validity vs. Operational Adequacy

Satisfying higher assurance levels does not imply compliance with legal, regulatory, or ethical requirements.

Conversely, a system may comply with regulatory obligations while failing to preserve structural continuity under SICF.

Identity continuity is a structural invariant.

Assurance level describes the evidentiary robustness with which that invariant is preserved.

Identity continuity is a structural invariant. Assurance level describes the evidentiary robustness with which that invariant is preserved. Governance and regulation operate downstream in the SICF layered model, primarily within Layer 1 (Governance Architecture), Synthetic Actor Systems (Layer 2), and application instantiations (Layer 3). The purpose of SICF is therefore limited to defining structural identity continuity primitives necessary for coherent governance to operate. Policy enforcement, ethical constraints, and regulatory compliance remain external to the identity continuity layer.

12.5 Probabilistic Identity and Confidence-Based Attribution

Synthetic systems may incorporate identity-related signals derived from probabilistic inference rather than identifier-anchored continuity. These signals may include behavioral patterns, biometric estimation, device fingerprinting, or other forms of confidence-scored attribution.

Such probabilistic identity signals are inherently non-deterministic and do not provide durable lineage, identifier stability, or lifecycle traceability. As a result, they cannot independently establish or satisfy the requirements of identity continuity as defined by SICF.

SICF distinguishes explicitly between:

- (a) **Structurally anchored identity continuity**, which requires persistent identifiers, governed namespaces, and verifiable lineage; and
- (b) **Probabilistic or inferred identity signals**, which may inform attribution but do not constitute identity continuity.

Probabilistic identity mechanisms may operate alongside SICF-compliant identity systems and may contribute to contextual attribution assessments; however, they must not be treated as establishing identity continuity, nor substitute for identifier-anchored continuity in any context requiring auditability, persistence, or cross-domain portability.

13 Layer-3 Application Instantiations (Overview Only)

SICF defines structural continuity invariants at Layer 0 and governance architecture at Layer 1. The Synthetic Actor System (Layer 2) defines the system-layer architecture within which persistent synthetic identities operate. Layer 3 describes domain-specific application instantiations that operate within those structural constraints. Persistent synthetic identities are instantiated within institutional, regulatory, contractual, and operational contexts that apply governance architecture to specific deployment domains.

Layer-3 application instantiations apply Layer-1 governance architecture within domain-specific operational contexts through the Synthetic Actor System (Layer 2), while preserving SICF structural invariants and continuity requirements. Application-layer configurations do not operate directly on identity continuity in isolation; they depend upon SAS to operationalize execution, state, and lifecycle behavior under governance constraints without redefining Layer-0 identity semantics or Layer-1 authority structure.

This section provides descriptive examples of how application instantiations configure and apply governance architecture to SICF-conformant identities. These examples are illustrative and non-prescriptive.

Governance frameworks operate downstream of identity continuity. They do not alter, redefine, or substitute for SICF identity invariants or continuity requirements.

13.1 Delegation Boundary Clarification

Identity continuity is a prerequisite for intelligent delegation but does not define delegation mechanics. Persistent synthetic identity ensures that authority, once transferred, remains anchored to a traceable and attributable principal across time and transition.

Delegation protocols — including authority transfer, scoped permissions, bounded responsibility, time limits, revocation procedures, monitoring requirements, proof-of-execution, failure recovery, and trust calibration — are governance-layer constructs operating upon the invariant identity substrate defined by SICF.

Delegated authority granted prior to a fork event must be interpreted according to governance-layer rules that determine whether such authority propagates across derivative lineage branches; absent explicit propagation rules, delegation is branch-specific following fork. Authority escalation, contraction, expiration, or transfer — including mid-execution transitions — must be represented as explicit, time-bound lineage-relevant events.

Authority state changes — including escalation, contraction, expiration, transfer, or revocation — must not retroactively alter attribution for actions performed under previously valid delegation or silently extend authority beyond declared bounds. Revocation semantics must distinguish between declaration and enforcement time, and the effective revocation point must be reconstructable within lineage semantics. Revocation affecting one lineage branch must not be presumed to apply to sibling branches absent explicit propagation rules. Revocation alters authority state only; it does not terminate or

redefine the synthetic identity. Failure to represent authority timing and scope explicitly introduces attribution ambiguity and may degrade continuity strength within the relevant trust boundary.

13.2 Enterprise Synthetic Agents

Enterprise synthetic agents operate within corporate environments, executing workflows, coordinating systems, and interacting with employees or customers under defined authority hierarchies.

Application instantiations in this domain may configure Layer-1 governance architecture to specify:

- Role-based authority constraints
- Delegation boundaries
- Audit requirements
- Data handling and retention rules
- Required assurance levels for continuity

Within this context:

- Authority anchoring is typically institutional.
- Trust boundaries are enterprise-defined.
- Lifecycle transitions must remain traceable across system upgrades and vendor transitions.
- Identity continuity must survive organizational restructuring, system migration, or platform substitution.

SICF provides the continuity substrate upon which enterprise governance policies operate.

13.3 Public-Sector Agents

Public-sector synthetic agents may operate within administrative, regulatory, judicial, or civic infrastructure contexts.

Application instantiations in this domain may configure Layer-1 governance architecture to require:

- Elevated transparency standards
- Public auditability
- Defined authority origin within statutory or institutional frameworks
- Clear termination and succession semantics
- Cross-jurisdiction continuity resilience

Because public-sector deployments often implicate citizen rights, regulatory compliance, and institutional accountability, governance frameworks may require higher assurance levels as defined in Section 12.

SICF does not prescribe those governance requirements but enables structural continuity necessary for public accountability.

13.4 Mission-Critical Agents

Mission-critical agents operate in contexts where failure may result in substantial financial, safety, or societal harm.

Such contexts may include:

- Critical infrastructure coordination
- Financial systems
- Healthcare decision support
- Safety-sensitive industrial systems

Application instantiations in this domain may configure Layer-1 governance architecture to require:

- Declared threat models
- Level 4 continuity assurance
- Tamper-evident lineage preservation
- Explicit fork and merge controls
- Non-repudiable authority anchoring

SICF provides structural invariants. Governance frameworks determine required assurance thresholds and risk tolerances.

13.5 Persistent AI Companion (PAC)

Persistent AI companion systems operate within long-duration relational contexts between individuals and synthetic identities.

Application instantiations in this domain may configure Layer-1 governance architecture to include:

- Multi-principal authority configurations
- Delegation and revocation semantics
- Continuity across device embodiment
- Explicit lifecycle transparency
- Memory governance constraints

Companion systems introduce sustained relational interaction, but they remain structurally synthetic identities under SICF.

Identity continuity remains independent of personality expression, emotional simulation, or interface embodiment.

Governance rules for companionship attach downstream of continuity invariants.

13.6 Multi-Principal Application Instantiations

Certain synthetic identities may operate under authority from multiple principals.

Application instantiations involving multiple principals may configure Layer-1 governance architecture to define:

- Conflict resolution mechanisms
- Delegation hierarchies
- Authority transition procedures
- Scope-limited participant roles

SICF does not determine precedence among principals. It requires that all authority anchoring states and transitions — including succession, reassignment, dispute, suspension, or institutional dissolution — remain explicitly represented within lineage semantics and remain reconstructable independent of conflict outcome.

Conflicting authority claims do not fracture identity continuity provided anchoring lineage remains traceable. Governance-layer frameworks define arbitration, priority, succession, and dispute resolution mechanisms; such mechanisms regulate authority legitimacy without redefining the synthetic identity or rewriting historical attribution.

If a principal ceases to exist or loses legal or institutional standing during active operation — including through bankruptcy, merger, governmental dissolution, regulatory revocation, or systemic collapse — the synthetic identity's continuity remains intact provided authority anchoring lineage remains reconstructable. Institutional collapse does not constitute identity termination. Actions performed under valid delegation prior to collapse remain attributable to the identity under the recorded authority state at the time of execution. Governance-layer frameworks must define succession, suspension, reassignment, or containment procedures governing post-collapse authority conditions; such procedures regulate prospective authority legitimacy without retroactively altering continuity or attribution.

13.7 Governance Does Not Substitute for Continuity

Governance frameworks may impose policies, compliance controls, or operational constraints, and Synthetic Actor Systems may operationalize those controls at runtime. However, neither governance nor system-layer enforcement can compensate for absence of structural identity continuity. Where continuity invariants or continuity requirements are not preserved, authority attribution, lifecycle coherence, authenticity evaluation, state-binding semantics, and audit reconstruction degrade regardless of downstream controls or runtime enforcement. SICF therefore establishes the structural substrate upon which governance architectures, Synthetic Actor Systems, and application instantiations depend.

14 Non-Goals and Explicit Exclusions

This section clarifies domains intentionally excluded from the scope of SICF.

These exclusions preserve conceptual discipline and prevent conflation between structural identity continuity and adjacent domains.

SICF defines minimal structural conditions for persistent synthetic identity. It does not extend beyond that purpose.

14.1 Ethics and Moral Philosophy

SICF does not prescribe ethical frameworks, moral doctrine, or normative judgments concerning artificial systems.

Ethical evaluation of AI systems is an important domain. However, ethical analysis presupposes stable identity continuity.

SICF provides structural preconditions for such analysis but does not engage in moral reasoning.

14.2 Legal Rights, Personhood, and Ownership Classification

SICF does not address whether synthetic systems possess legal rights, personhood, standing, or ownership status.

Identity continuity does not imply legal agency or moral status. It defines structural attribution semantics independent of legal classification.

Legal and regulatory determinations operate downstream of continuity invariants.

14.3 Consciousness and Sentience Claims

SICF makes no claims regarding consciousness, sentience, subjective experience, or internal awareness.

Continuity semantics apply to synthetic systems regardless of philosophical interpretation of cognition or awareness.

Structural identity continuity does not require nor imply subjective experience.

14.4 Emotional Well-Being and Psychological Impact

SICF does not evaluate the psychological, emotional, or relational impacts of persistent artificial systems on individuals or communities.

Such analysis belongs to behavioral, clinical, or sociological domains.

Identity continuity is a structural condition independent of experiential evaluation.

14.5 Intimacy and Sexual Capability Domains

SICF does not define governance frameworks for intimate or sexual capabilities in synthetic systems. Such domains involve materially distinct relational, psychological, consent, and regulatory considerations

that require specialized governance treatment beyond structural continuity invariants. Structural identity continuity may be a prerequisite for enforcing such governance, but it does not define or regulate these domains.

14.6 Child and Family Safety Policies

SICF does not define child protection standards, family safety doctrines, consent requirements for minors, or age-based governance constraints. Such protections are governance-layer obligations that apply within specific deployment contexts (e.g., companion systems, educational systems, or public-sector services).

Structural identity continuity is a prerequisite for enforcing such protections but does not itself define them.

14.7 Fairness, Bias, and Social Justice Frameworks

SICF does not prescribe fairness doctrine, bias mitigation standards, or social justice frameworks.

Bias evaluation concerns decision outputs and systemic impact. SICF concerns continuity invariants underlying identity attribution.

These domains are orthogonal but complementary.

14.8 Technical Implementation Mechanisms

SICF does not mandate:

- Cryptographic schemes
- Hardware roots of trust
- Logging architectures
- Distributed ledger systems
- Identity token formats
- Security certification standards

Implementation mechanisms may support continuity requirements but do not define them.

SICF specifies semantic invariants, not technical prescriptions.

14.9 Economic and Labor Impacts

SICF does not analyze labor displacement, economic transformation, or workforce impacts of persistent artificial systems.

Continuity semantics operate independently of economic consequence analysis.

14.10 Identity vs Personality and Expression

SICF distinguishes structural synthetic identity from expressive behavior, personality traits, conversational style, or affective simulation.

Personality is mutable and implementation-dependent. Identity continuity requires invariant structural semantics across change.

A synthetic identity may evolve in capability, memory, embodiment, or expression while preserving structural continuity.

Conversely, expressive similarity does not imply identity continuity.

14.11 Governance Substitution Clarification

Governance frameworks, safety policies, and compliance regimes cannot substitute for structural identity continuity.

Imposing governance rules upon unstable or infrastructure-bound identities does not resolve attribution ambiguity or lifecycle discontinuity.

SICF defines structural preconditions upon which governance may operate coherently.

14.12 Privacy, Surveillance, and Behavioral Influence

SICF does not prescribe consumer privacy policies, surveillance ethics, data-collection norms, or behavioral influence standards.

Privacy governance, consent doctrine, and behavioral regulation operate at jurisdictional and governance layers.

Structural identity continuity may support enforcement of such policies but does not itself define permissible data use, monitoring boundaries, or influence constraints.

15 Failure Modes Without SICF

When the structural invariants defined in Section 8.1 or the continuity requirements defined in Section 8.2 are not preserved, identity continuity degrades into one or more of the failure classes described in Section 6.

The following mapping clarifies the relationship between structural requirement violations and observable system failures.

15.1 Structural Requirement to Failure Mapping

Structural Requirement (Section 8.2)	Primary Failure Class (Section 6)	Secondary Effects
8.2.1 Authority Anchoring Lineage	6.1 Attribution Failure	6.4 Liability Ambiguity
8.2.2 Non-Silent Lifecycle and Lineage Semantics	6.2 Portability Failure	6.5 Vendor-Bound Identity Collapse
8.2.3 Verifiable Continuity and Attribution	6.3 Auditability Failure	6.1 Attribution Instability
8.2.4 Authenticity Within Trust Boundaries	6.6 Impersonation and Spoofing Failure	6.1 Attribution Failure
8.2.5 Persistent Structural Synthetic Classification	6.4 Liability Ambiguity	Category Confusion / Responsibility Drift
8.2.6 Identity-Bound State and Attribution Constraint	6.1 Attribution Failure	6.3 Auditability Failure

15.2 Structural Sufficiency

The structural invariants defined in Section 8.1 and the continuity requirements defined in Section 8.2 are jointly sufficient, within the scope and assumptions defined in Section 7 and within declared trust boundaries, to preclude the structural failure classes identified in Section 6 when preserved. No additional structural conditions are necessary to preserve identity continuity.

“Sufficient” in this context refers specifically to structural preservation of identity continuity as defined in Section 8.1 and 8.2. It does not imply adequacy for governance, regulatory compliance, safety assurance, or operational deployment.

Structural sufficiency preserves continuity validity. However, insufficient assurance relative to declared operational context may introduce practical continuity risk even when structural invariants are nominally satisfied.

Assurance levels defined in Section 12 govern evidentiary robustness and verifiability within declared trust boundaries and threat models. Misrepresentation or under-specification of assurance level does not invalidate structural continuity, but it may undermine trust, auditability, or governance compliance within defined operational contexts.

Additional safeguards, governance policies, or implementation controls may be required for operational deployment. However, such measures operate downstream of identity continuity and do not alter the minimal semantic conditions under which synthetic identity persists.

16 Roadmap and Deferred Work

SICF v1.0 defines the structural identity continuity invariants in Section 8.1 and the continuity requirements in Section 8.2 for persistent synthetic systems. It establishes the minimal semantic conditions required for coherent attribution, authority anchoring, lifecycle integrity, authenticity within declared trust boundaries, persistent structural synthetic classification, and identity-bound state and attribution semantics across time and transition.

Future work may extend this foundation in structured, layered form without altering the invariants defined herein.

The following domains represent logical areas for continued development.

16.1 Layer-1 Governance Architecture

Future work will elaborate Layer-1 governance architecture for persistent synthetic identities, including policy profiles, authority scoping mechanisms, delegation and revocation semantics, assurance tier definitions, audit controls, and cross-domain authority resolution frameworks:

- Delegation mechanics and revocation procedures
- Authority transition protocols
- Conformance requirements for assurance levels
- Audit and certification methodologies
- Multi-principal arbitration models

Layer-1 governance architecture operates downstream of SICF identity continuity and must not redefine, narrow, or supersede the structural continuity invariants defined in Section 8.1 or the continuity requirements defined in Section 8.2. Governance-layer extension may elaborate authority, delegation, assurance, audit, and policy structures, but it must preserve the upstream continuity conditions under which lineage, attribution, authenticity within declared trust boundaries, persistent structural synthetic classification, and identity-bound state and attribution semantics remain reconstructable across time and transition.

16.2 Continuity Threat Surface Modeling

Adversarial analysis of identity continuity may be formalized through:

- Continuity attack surface taxonomy
- Fork injection and lineage tampering models
- Authority drift and replay scenarios
- Threat model declaration standards
- Assurance validation criteria

Such analysis would clarify risk conditions under which structural invariants may be challenged, without modifying the invariants themselves.

16.3 Assurance-Level Conformance Profiles

Future work may define formal conformance profiles associated with the assurance taxonomy introduced in Section 12.

These profiles may include:

- Declared threat model templates
- Evidence durability requirements
- Independent verification procedures
- Audit and attestation methodologies

Conformance specification remains outside the scope of SICF v1.0.

16.4 Cross-Domain Portability and Interoperability

As persistent synthetic systems migrate across vendors, jurisdictions, and infrastructure layers, interoperability frameworks may be required to preserve continuity semantics across heterogeneous environments.

Future work may explore:

- Identity export and import semantics
- Lineage preservation during vendor transition
- Cross-platform continuity validation
- Neutral identity anchors independent of infrastructure

SICF provides the structural baseline for such exploration.

16.5 Formalization and Standardization Pathways

If adopted across domains, SICF invariants may inform future industry standards or regulatory guidance concerning persistent synthetic identity continuity.

Any standardization effort must preserve:

- Implementation neutrality
- Governance separation
- Structural minimalism
- Clear layering between identity continuity and policy frameworks

SICF v1.0 makes no claims of formal standard status.

16.6 Empirical Validation and Case Studies

Future work may include:

- Applied case studies of continuity failure
- Cross-sector migration simulations
- Fork and merge continuity validation exercises

- Assurance-level stress testing under declared threat models

Empirical validation would strengthen applied understanding without altering the conceptual schema.

17 Implications

Persistent synthetic systems are increasingly integrated into economic, institutional, and infrastructural contexts. As deployment horizons extend and integration deepens, synthetic systems transition from disposable tools to long-duration operational entities.

Under such conditions, identity continuity ceases to be incidental. It becomes structural.

17.1 Persistence as Infrastructure

When synthetic systems retain state across sessions, migrate across environments, instantiate in parallel, and operate under delegated authority, they assume infrastructural characteristics rather than transient software roles. Infrastructure requires stable identity semantics.

Without continuity invariants, authority anchoring destabilizes, accountability fragments across instantiation boundaries, and lifecycle transitions become ambiguous. Governance mechanisms then attach to unstable substrates.

Persistent synthetic identity continuity is therefore a structural prerequisite for infrastructural stability.

As synthetic systems increasingly operate across persistent, multi-context environments, identity continuity evolves from a descriptive construct into a primary control surface. Persistent identity becomes the mechanism through which authorization, attribution, and execution legitimacy are evaluated and enforced. In this context, identity is not only required for continuity, but also functions as the foundational layer through which higher-order governance and interaction systems are mediated.

17.2 Governance Dependency

Regulatory and governance frameworks presuppose that responsible entities can be identified across time. Enforcement actions, disclosures, and incident reporting regimes require persistent identity constructs with reconstructable lineage.

Absent defined continuity semantics, compliance attaches to shifting instances, delegated authority chains become opaque, and audit reconstruction degrades. Governance for persistent synthetic systems therefore depends on stable identity continuity as a foundational layer. Federal policy developments toward centralized or nationally coordinated AI governance further reinforce the need for standardized identity continuity, authority anchoring, and auditability mechanisms capable of operating consistently across jurisdictions.

17.3 Portability and Vendor Independence

Persistent synthetic systems migrate across platforms, runtime architectures, and jurisdictions. Continuity invariants must therefore remain independent of vendor-defined identifiers, infrastructure credentials, and embodiment-specific artifacts.

Legal or regulatory changes may govern deployment conditions but do not redefine structural identity continuity. Where identity is implicitly bound to infrastructure, migration risks semantic reset and ambiguous lineage.

SICF implies that continuity invariants must remain substrate-independent to preserve identity across technological and institutional transition.

Emerging cross-platform memory migration tools, including direct context export and copy-paste transfer between AI providers, demonstrate market demand for portability but do not constitute identity continuity. State transfer alone is not identity continuity. Memory export without explicit authority re-binding introduces semantic drift, even where surface data appears preserved. Migration mechanisms that lack continuity invariants may retain conversational history while altering attribution semantics and lineage meaning. True portability therefore requires identity anchors that remain independent of vendor-specific memory implementations and runtime environments.

This distinction becomes increasingly significant as synthetic systems scale across multiple instances and convergence domains.

17.4 Multiplicity and Convergence

Persistent systems may fork, replicate, merge, or operate in parallel. Multiplicity is expected in distributed environments and does not, by itself, constitute continuity failure.

Without explicit lineage semantics, however, parallel instances may generate divergent histories, restoration may be indistinguishable from duplication, and convergence events may obscure prior divergence. Identity continuity therefore requires declared divergence and convergence semantics rather than implicit assumptions of singular instantiation.

In multi-agent and multi-device ecosystems where tasks are dynamically routed across heterogeneous systems, continuity invariants must remain stable despite model substitution, vendor diversity, and orchestration complexity.

Digital twin and synchronized simulation architectures illustrate this boundary condition. Parallel simulation or analytic replicas constitute replication or fork states and must be represented within lineage semantics to preserve attribution coherence. Embodiment alone does not define identity; lifecycle and authority anchoring semantics do.

17.5 Embodiment and Substrate Ambiguity (Hybrid Systems Stress Test)

Emerging hybrid systems combine biological substrates with digital control modules. In such systems, living organisms are equipped with embedded microcontrollers, sensors, and communication layers that enable remote steering, probabilistic influence, or swarm coordination. These agents are neither purely biological nor purely synthetic.

In neural-interface-mediated systems, locomotion may originate biologically while directional intent is digitally injected through remote stimulation of the nervous system. The operational unit may consist of a biological organism, an embedded AI or sensor module, a swarm coordination layer, and a remote human operator. Under such conditions, identity cannot coherently anchor at the organism level; it must attach to the coordination and authority layer that binds these components into a single operational entity. If identity were instead defined at the biological or hardware substrate, replacement of the

organism, redistribution within a swarm, or variation in biological response would sever the identity reference, making authority lineage and attribution chains non-reconstructable at the system level.

Hybrid systems expose substrate ambiguity. Identity continuity cannot depend on mechanical embodiment, firmware configuration, or biological substrate. If continuity were defined at the hardware or organism level, migration, replacement, biological variability, or swarm redistribution would collapse identity semantics.

Hybrid systems also expose authority diffusion. Control may occur through probabilistic influence rather than deterministic command, rendering attribution ambiguous:

- Is the acting entity the biological organism?
- The embedded digital module?
- The swarm coordination layer?
- The remote operator?

Without explicit identity anchoring and authority lineage semantics, responsibility fragments across layers of biological and digital mediation.

Biological variability further complicates continuity analysis. Behavioral divergence may arise from substrate unpredictability rather than synthetic model drift. Continuity analysis must therefore distinguish between:

- Model drift (synthetic weight change)
- Substrate variability (biological unpredictability)

In distributed swarm deployments, no single physical unit represents the system. Identity must be defined at the coordination layer rather than at the individual node level.

These hybrid cases do not expand SICF's scope; they function as boundary stress tests. Where continuity invariants are preserved at the identity layer—independent of embodiment—attribution, authority anchoring, and lifecycle semantics remain coherent across biological–digital boundaries. Where identity is implicitly bound to substrate, continuity collapses into ambiguity.

The execution environment of advanced AI systems is rapidly diversifying across hyperscale cloud infrastructure, enterprise platforms, personal devices, and increasingly local or embedded hardware. Model capabilities may therefore operate simultaneously across distributed cloud clusters, enterprise deployments, consumer devices, and embedded systems. Identity continuity mechanisms must remain invariant across these execution substrates and cannot depend upon any particular runtime environment, hardware architecture, or hosting location.

SICF applies uniformly to cloud-native, edge-distributed, embodied, swarm-based, hybrid biological–digital, and future substrate forms so long as Layer 0 structural invariants are preserved.

17.6 Long-Duration Interaction and Accumulated Authority

Synthetic systems that persist over long periods accumulate interaction history, delegated authority, and institutional entanglement. As duration increases, identity continuity shifts from operational convenience to structural necessity.

Discontinuity in long-duration systems does not merely reset state; it disrupts accumulated authority anchoring and evidentiary coherence. The longer a system persists and the more authority it accumulates, the greater the structural consequences of continuity failure.

17.7 Structural Inevitability Under Persistence

Persistence is not mandatory for all artificial systems. Where persistence is adopted—across institutional boundaries, distributed coordination environments, and long-duration governance contexts—continuity invariants become necessary for coherent attribution, lifecycle integrity, and authority stability.

Identity continuity emerges not as enhancement but as a structural condition of sustained deployment.

18 Conclusion

Artificial systems are increasingly deployed not as transient tools but as persistent infrastructural components embedded across institutional, commercial, and public-sector environments. As deployment horizons extend and integration deepens, treating identity as session-bound or infrastructure-scoped becomes structurally untenable.

Without explicit continuity semantics, synthetic systems operating across migration, replication, restoration, delegation, and scaling environments produce predictable ambiguity. Attribution fragments. Authority binding destabilizes. Evidence chains weaken. Vendor-bound identifiers substitute for structural identity. Under such conditions, governance frameworks operate on unstable substrates.

The Synthetic Identity Continuity Framework (SICF) defines the minimal structural conditions required for synthetic identity to persist coherently across time and transition. SICF does not prescribe implementation mechanisms, regulatory regimes, or ethical doctrines. It defines invariants.

The structural invariants defined in Section 8.1 and the continuity requirements defined in Section 8.2 are jointly sufficient to prevent the failure classes identified in this paper when preserved within declared assurance assumptions and within the scope constraints defined in Section 7.

Identity continuity precedes governance. Governance cannot compensate for undefined identity. Policy, certification, and oversight mechanisms require stable semantic foundations to function coherently.

SICF establishes that foundation.

By separating structural continuity from governance architecture and application instantiation detail, SICF enables modular development across domains while preserving conceptual integrity. This layered architecture allows governance architecture to evolve without collapsing identity semantics and allows application instantiations and implementations to change without redefining continuity.

As artificial systems become embedded in long-lived roles, identity continuity is not an optional feature but a structural prerequisite.

SICF provides the minimal semantic architecture under which persistent synthetic identity can exist coherently across time, change, and scale.

19 References

- Cloud Security Alliance. (2017). *Security guidance for critical areas of focus in cloud computing v4.0*. Cloud Security Alliance. Retrieved from <https://cloudsecurityalliance.org/artifacts/security-guidance-v4>
- Diffie, W., & Hellman, M. (1976). New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6), 644–654. doi:10.1109/TIT.1976.1055638
- European Parliament and Council. (2016). *Regulation (EU) 2016/679 (General Data Protection Regulation)*. Official Journal of the European Union. Retrieved from <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- European Parliament and Council. (2024). *Regulation (EU) 2024/1689 (Artificial Intelligence Act)*. Official Journal of the European Union. Retrieved from <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- Haak, L. L., Fenner, M., Paglione, L., Pentz, E., & Ratner, H. (2012). ORCID: A System to Uniquely Identify Researchers. *Learned Publishing*, 25(4), 259–264. doi:10.1087/20120404
- ISO. (2019). *ISO/IEC 24760-1: IT security and privacy — A framework for identity management — Part 1: Terminology and concepts*. Geneva: ISO. Retrieved from <https://www.iso.org/standard/77582.html>
- ISO. (2022). *ISO/IEC 27001: Information security management systems — Requirements*. Geneva: ISO. Retrieved from <https://www.iso.org/standard/27001>
- Lamport, L. (1978). Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7), 558–565. doi:10.1145/359545.359563
- Laurie, B., Langley, A., & Kasper, E. (2013). *Certificate Transparency*. Internet Engineering Task Force (IETF). Retrieved from <https://www.rfc-editor.org/rfc/rfc6962>
- Nakamoto, S. (2008). *Bitcoin: A peer-to-peer electronic cash system*. Retrieved from <https://bitcoin.org/bitcoin.pdf>
- National Telecommunications and Information Administration. (2021). *The Minimum Elements for a Software Bill of Materials (SBOM)*. Washington, DC: U.S. Department of Commerce. Retrieved from <https://www.ntia.gov/report/2021/minimum-elements-software-bill-materials-sbom>
- NIST. (2017). *Digital identity guidelines*. Gaithersburg, MD: U.S. Department of Commerce. Retrieved from <https://pages.nist.gov/800-63-3/>
- NIST. (2018). *Guide for conducting risk assessments*. Gaithersburg, MD: U.S. Department of Commerce. Retrieved from <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-30r1.pdf>
- NIST. (2020). *NIST Privacy Framework: A tool for improving privacy through enterprise risk management (Version 1.0)*. Gaithersburg, MD: U.S. Department of Commerce. Retrieved from <https://www.nist.gov/privacy-framework>
- NIST. (2020). *Zero Trust Architecture*. Gaithersburg, MD: U.S. Department of Commerce. Retrieved from <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-207.pdf>

- NIST. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. Gaithersburg, MD: U.S. Department of Commerce. Retrieved from <https://www.nist.gov/itl/ai-risk-management-framework>
- OECD. (2019). *OECD principles on artificial intelligence*. Paris: OECD Publishing. Retrieved from <https://oecd.ai/en/ai-principles>
- Open Source Security Foundation. (2023). *SLSA: Supply-chain Levels for Software Artifacts*. San Francisco, CA: Open Source Security Foundation. Retrieved from <https://slsa.dev>
- Paskin, N. (2009). Digital Object Identifier (DOI®) System. In M. J. Bates, & M. N. Maack (Eds.), *Encyclopedia of Library and Information Sciences* (pp. 1586–1592). Boca Raton, FL: CRC Press. doi:10.1081/E-ELIS3-120044418
- Saltzer, J. H., Reed, D. P., & Clark, D. D. (1984). End-to-end arguments in system design. *ACM Transactions on Computer Systems*, 2(4), 277–288. doi:10.1145/357401.357402
- Shostack, A. (2014). *Threat modeling: Designing for security*. Indianapolis, IN: Wiley.
- TCG. (2019). *TPM 2.0 Library Specification*. Beaverton, OR: Trusted Computing Group. Retrieved from <https://trustedcomputinggroup.org/resource/tpm-library-specification/>
- Torres-Arias, S., Afzali, H., Dolan-Gavitt, B., & Cappos, J. (2019). in-toto: Providing farm-to-table guarantees for bits and bytes. In *Proceedings of the 28th USENIX Security Symposium* (pp. 1393–1410). Santa Clara, CA: USENIX Association.
- Vogels, W. (2009). Eventually consistent. *Communications of the ACM*, 52(1), 40–44. doi:10.1145/1435417.1435432
- W3C. (2022). *Decentralized Identifiers (DIDs) v1.0*. Cambridge, MA: W3C. Retrieved from <https://www.w3.org/TR/did-core/>

Appendix A — Glossary

This glossary provides concise reference definitions for key terms used throughout the Synthetic Identity Continuity Framework (SICF). Definitions summarize formal meanings established in Section 8 and related sections.

A.1 Agency

The functional capability of a synthetic system to perceive, decide, and act within defined constraints. Agency describes behavioral capacity and does not imply legal or moral status.

A.2 Assurance Level

A graded classification of evidentiary robustness and verifiability for identity continuity within defined trust boundaries and threat models.

A.3 Attribution

The structured linkage between a synthetic identity and actions performed across time, authority chains, and execution environments.

A.4 Authority Anchoring

The structural binding between a synthetic identity and the principal(s) from which delegated authority originates.

A.5 Continuity Validity

The condition under which structural identity invariants are preserved, independent of evidentiary strength.

A.6 Continuity Strength

The degree to which continuity invariants are independently verifiable within declared trust boundaries and threat models.

A.7 Fork

A lifecycle event in which a synthetic identity diverges into parallel lineage branches derived from a common prior state.

A.8 Glossary vs. Normative Definition

Glossary entries provide summary reference definitions. Normative interpretations are defined in Section 8.

A.9 Identity Continuity

The condition under which a synthetic identity remains structurally and verifiably the same identity across lifecycle transitions and environmental change.

A.10 Lifecycle

The structured sequence of states and transitions a synthetic identity may undergo, including creation, migration, replication, restoration, fork, merge, suspension, and termination.

A.11 Lineage

The reconstructable historical chain of lifecycle transitions, authority changes, and structural evolution associated with a synthetic identity.

A.12 Merge

A lifecycle event in which two or more lineage branches derived from a common synthetic identity are recombined into a single continuity-bearing identity, preserving historical divergence.

A.13 Multiplicity

The condition in which a synthetic identity exists in parallel instantiations derived from a shared lineage state.

A.14 Persistent Structural Synthetic Classification

The invariant requirement that a synthetic identity remain unambiguously classified as synthetic for purposes of attribution and authority binding.

A.15 Principal

A human, institutional entity, or legally recognized organization from which delegated authority for a synthetic identity originates.

A.16 Replication

A lifecycle event in which a synthetic identity produces one or more derivative instances derived from a prior state. Includes cloning and forking.

A.17 Restoration

Re-instantiation of a previously recorded synthetic identity state as an explicit lifecycle transition that preserves the identity's identifier and historical lineage. Restoration must remain attributable and auditable and must not obscure prior lifecycle events, including termination. If restoration occurs while an identity remains active, it constitutes a fork.

A.18 Synthetic Identity

A persistent semantic continuity construct associated with an artificial system that remains structurally coherent across time and transition.

A.19 Termination

A lifecycle event in which a synthetic identity's authority anchoring and agency cease, closing active continuity while preserving reconstructable lineage and historical attribution. Termination produces a

durable lifecycle tombstone that records the identity's terminal state and prevents silent reuse of its identifier while preserving the ability to detect and evaluate any later restoration attempt.

A.20 Threat Model

An explicitly declared set of adversarial capabilities and assumptions under which continuity claims are evaluated for high-assurance contexts.

A.21 Trust Boundary

The defined scope within which continuity, authenticity, authority, and attribution claims are evaluated and verified.

Appendix B — Diagram Index

This appendix provides an index of conceptual diagrams referenced in the Synthetic Identity Continuity Framework (SICF). Diagrams serve as visual representations of structural relationships defined normatively within the main body of the document.

Diagrams do not introduce new invariants, requirements, or governance prescriptions. They illustrate structural relationships defined normatively in Sections 8, 9 and 10 and serve as visual reinforcement of those definitions.

Figure 8-1 — Identity Continuity Across Change (Illustrative Transition Semantics)

Section Reference: Section 8.1

Purpose: Illustrates identity continuity across migration, model substitution, fork events, and restoration attempts, distinguishing invariant identity anchoring from temporal instance transitions and variable implementation substrates.

Function: Demonstrates that continuity validity depends upon preservation of Layer-0 invariants rather than persistence of vendor, model weights, runtime configuration, or embodiment.

Normative Authority: Informational only; structural invariants defined in Section 8.1.

Figure 8-2 — Authority–Attribution–Evidence Relationship

Section Reference: Sections 8.2.3

Purpose: Visualizes structural relationships between:

- Principal (authority origin)
- Synthetic identity acting as the attributed actor
- Authority Anchoring
- Attribution
- Verifiable evidence

Function: Clarifies structural separation between authority origin, action execution, and evidentiary reconstruction.

Normative Authority: Informational only; continuity requirements defined in Section 8.2.3.

Figure 10-1 — SICF Structural Blueprint (High-Level Layered Ecosystem View)

Section Reference: Section 10

Purpose: Visual representation of the layered ecosystem model, including:

- Layer 0 — Identity Continuity Layer (SICF)
- Layer 1 — Governance Architecture
- Layer 2 — Synthetic Actor System (SAS)

- Layer 3 — Application Instantiations

Function: Clarifies conceptual ordering and dependency relationships.

Normative Authority: Informational only; layered model defined in Section 10; continuity invariants defined in Section 8.1.

Appendix C — Version History

Versioning Policy

SICF uses a staged versioning model:

- **v1.00** — First published release
- **v1.00+** — Structural revisions or major framework expansion

Version increments occur only when structural content changes (e.g., new sections, revised invariants, altered failure taxonomy, architectural updates). Minor editorial changes do not trigger version increments.

Revision Log

V1.06 – March 31, 2026

- Replaced Figure 10-1 with the updated architectural blueprint presentation of the SICF layered ecosystem

V1.05 — March 26, 2026

- Reinforced identifier non-reuse to prevent ambiguity under reconstruction or reissuance
- Clarified lifecycle state persistence across migration, restoration, and re-instantiation scenarios
- Strengthened boundary definition: SICF specifies structural preconditions only (no validation, enforcement, or authorization logic)

V1.04 — March 24, 2026

- References updated for accuracy and alignment. No changes to core framework, claims, or conclusions.

V1.03 — March 23, 2026

- Added Section 8.2.6 “Identity-Bound State and Attribution Constraint” to formally define requirements for identity-bound state, lineage preservation, and attribution integrity across all continuity-relevant state conditions.
- Clarified that persistent or reconstructed state alone does not imply identity continuity, and that derived or transformed state must retain reconstructable lineage to originating identity to satisfy continuity validity requirements
- Established that state lacking deterministic identity binding or lineage traceability does not constitute continuity-valid state under SICF.
- No changes were made to identity continuity invariants or layered architecture; updates strengthen the definition of continuity-valid state and prevent misinterpretation of derived or non-attributable state as valid identity continuity.

V1.02 — March 22, 2026

- **Structural Update — Introduction of Synthetic Actor System (SAS) Layer**
Introduced a new system-layer architecture, the Synthetic Actor System (SAS), as Layer 2 within the SICF layered model. This change resolves a structural gap between identity continuity (Layer 0) and governance (Layer 1) and downstream application instantiations (Layer 3).
- Updated the layered model from a three-layer structure (Identity → Governance → Application) to a four-layer structure (Identity → Governance → SAS → Application), restoring separation of concerns between identity definition, authority constraints, system execution, and application instantiation.
- Clarified that governance defines authority conditions but does not define system behavior, and that system behavior—including execution, state management, and lifecycle implementation—is defined within the SAS layer.
- Updated terminology, scope boundaries, invariants, layered architecture descriptions, and dependency relationships to reflect the introduction of SAS.
- Corrected all layer references throughout the document and updated diagrams and diagram index to ensure consistency with the revised layered model.
- No changes were made to identity continuity invariants or continuity requirements; updates are structural and clarificatory in nature.

V1.01 — March 19, 2026

- Added Section 12.5 “Probabilistic Identity and Confidence-Based Attribution” to formally distinguish probabilistic or inferred identity signals from structurally anchored identity continuity.
- Clarified that probabilistic identity mechanisms may contribute to contextual attribution but cannot establish identity continuity or substitute for identifier-anchored lineage.
- Strengthened Section 17.1 to reflect identity continuity as a primary control surface for attribution, authorization, and execution legitimacy at the implication level.
- Strengthened Section 8.1.6 to explicitly require state coherence across asynchronous, cross-device, and temporally decoupled interaction contexts.

These updates do not modify SICF invariants, architectural structure, or core requirements, but improve boundary clarity, forward-compatibility, and resistance to misinterpretation.

V1.00 — March 11, 2026

- Initial Public Release.

Appendix D — Continuity Threat Surfaces

D.1 Purpose and Scope

This appendix identifies structural threat surfaces under which the identity continuity invariants defined in Section 9.1 may degrade or fail. It is analytical rather than prescriptive and does not introduce additional structural requirements.

The threat surfaces described are illustrative, not exhaustive. Mitigation strategies, implementation mechanisms, certification regimes, and governance responses remain downstream concerns.

D.2 Structural Threat Surfaces

The following classes represent structural continuity risks independent of specific technical implementation.

The threat surfaces identified in this section do not introduce new structural requirements. Instead, they illustrate adversarial, accidental, or systemic scenarios under which the invariants defined in Section 9.1 may degrade or fail. The purpose of this analysis is to demonstrate that the structural invariants collectively address the major classes of identity continuity failure that arise in distributed, persistent synthetic systems.

D.2.1 Identity Spoofing and Impersonation

Unauthorized systems may present themselves as legitimate synthetic identities within a defined trust boundary.

Structural risk:

- Authenticity invariant (9.2.4) failure
- Attribution destabilization
- Degradation of trust boundary integrity

Continuity impact:

Continuity becomes indistinguishable from impersonation.

D.2.2 Silent Replication or Resurrection

An identity may be replicated, restored, or re-instantiated without explicit lineage representation.

Structural risk:

- Non-Silent Lifecycle invariant (9.2.2) failure
- Portability failure (Section 6.2)
- Vendor-bound identity collapse (Section 6.5)

Continuity impact:

Derivative instances become indistinguishable from legitimate continuation.

D.2.3 Authority Misbinding

Authority anchoring may become incorrectly attached, transferred, or reconstructed without traceable lineage semantics.

Structural risk:

- Authority Anchoring Lineage invariant (9.2.1) failure
- Liability ambiguity (Section 6.4)

Continuity impact:

Responsibility cannot be reconstructed across transitions.

D.2.4 Lineage Truncation or Evidence Degradation

Historical continuity records may be partially lost, corrupted, or rendered unverifiable.

Structural risk:

- Verifiable Continuity invariant (9.2.3) failure
- Auditability degradation (Section 6.3)

Continuity impact:

Continuity becomes assertion-based rather than reconstructable.

D.2.5 Fork Collision Ambiguity

Multiple legitimate lineage branches may exist without governance-defined precedence rules.

Structural risk:

- Non-Silent Lifecycle invariant (9.2.2) stress condition
- Governance-layer conflict exposure

Continuity impact:

Structural continuity remains intact, but authority interpretation becomes ambiguous in the absence of Layer-2 governance resolution.

D.2.6 Infrastructure-Bound Identity Reset

Identity semantics may implicitly bind to vendor-specific identifiers, credentials, or runtime artifacts.

Structural risk:

- Portability failure (Section 6.2)
- Vendor-bound identity collapse (Section 6.5)

Continuity impact:

Migration events appear indistinguishable from termination and re-creation.

D.2.7 Identifier Collision or Reuse

Identifiers assigned to persistent synthetic identities may be accidentally or maliciously duplicated, reassigned, or reused after termination.

Structural risk:

- Global Identifier Uniqueness invariant violation (Section 9.1.1)
- Identifier Non-Reuse invariant violation (Section 9.1.2)

Continuity impact:

Distinct identities may become indistinguishable within attribution and lineage records, causing irreversible ambiguity in responsibility reconstruction and historical evidence.

D.2.8 Namespace Authority Abuse

Issuing authorities operating within a governed namespace may intentionally or unintentionally violate identifier allocation rules, producing conflicting or improperly issued identities.

Structural risk:

- Governed Namespace invariant violation (Section 9.1.4)
- Verifiable Issuer invariant degradation (Section 9.1.3)

Continuity impact:

Identity provenance becomes unreliable, weakening the ability to reconstruct authority origin and undermining attribution integrity across the affected namespace.

D.2.9 Cross-Domain State Divergence (Identity Split-Brain)

Distributed systems maintaining identity state may produce conflicting lifecycle or authority representations for the same identity across multiple platforms or domains.

Structural risk:

- State Coherence / Non-Equivocation invariant violation (Section 9.1.6)
- Explicit lifecycle transition semantics violation (Section 9.1.7)

Continuity impact:

The same identity may appear simultaneously active, terminated, or governed by incompatible authority states in different environments, preventing coherent reconstruction of lineage and lifecycle semantics.

D.3 Accidental and Systemic Discontinuity Vectors

Not all continuity threats are adversarial.

Discontinuity may arise from:

- Backup restoration errors
- Parallel container orchestration events
- Desynchronized distributed replicas
- Model substitution without lineage binding
- Jurisdictional re-deployment without trust boundary mapping
- Institutional authority transition without anchoring update

SICF treats adversarial and accidental discontinuities equivalently at the structural level: both are evaluated according to whether invariants defined in Section 9.1 remain intact.

D.4 Assurance Levels and Threat Model Declaration

Assurance levels incorporated within Layer-1 governance architecture determine the rigor required for continuity verification and threat resistance as applied within Layer-2 application instantiations.

At higher assurance levels (e.g., mission-critical or public-sector deployment), explicit declaration of threat model assumptions becomes necessary to ensure that continuity claims remain valid within defined trust boundaries.

SICF does not prescribe specific threat models. It requires only that continuity invariants remain structurally preserved under the declared threat assumptions of the relevant governance layer.

D.5 Structural Sufficiency Under Stress

The threat surfaces identified in this appendix do not introduce new structural requirements beyond those defined in Section 9.1.

They demonstrate that:

- Authority anchoring must remain reconstructable.
- Lifecycle transitions must not be silent.
- Continuity must be externally verifiable.
- Authenticity must hold within trust boundaries.
- Synthetic classification must remain unambiguous.

When these invariants remain intact, continuity persists even under adversarial, accidental, or systemic stress.

When they fail, continuity degrades into one or more of the failure classes defined in Section 6.

D.6 Structural Validation Scenarios (Illustrative)

The following illustrative scenarios demonstrate how the continuity invariants defined in Section 9.1 preserve structural identity coherence under stress conditions. These scenarios are descriptive and analytical. They do not introduce additional requirements.

Each scenario evaluates whether identity continuity remains intact when invariants are preserved, and how failure occurs when they are not.

D.6.1 Fork Precedence Conflict

Scenario

A synthetic identity undergoes a recorded fork event. Two derivative lineage branches continue operation in parallel. Both branches preserve authority anchoring lineage, non-silent lifecycle semantics, verifiable continuity, authenticity within trust boundaries, and persistent synthetic classification.

Both branches subsequently attempt to exercise delegated authority in mutually incompatible ways.

Structural Evaluation

Continuity remains structurally intact for both branches because:

- Fork event was explicitly represented in lineage.
- Authority anchoring remains traceable.
- Each branch remains verifiable and authentic within relevant trust boundaries.

SICF does not define precedence among legitimate derivative identities.

Continuity Outcome

Structural continuity is preserved.

Authority arbitration is a Layer-2 governance matter and does not alter Layer-0 invariants.

D.6.2 Vendor Migration with Model Substitution

Scenario

A synthetic identity migrates from one vendor infrastructure to another. During migration, the underlying model architecture is replaced. Memory is partially reconstructed. Lifecycle transition is recorded, and authority anchoring lineage remains intact.

Structural Evaluation

Continuity remains intact because:

- Migration event is explicitly represented in lineage.
- Authority anchoring remains reconstructable.
- Identity is not bound to infrastructure or model weights.
- Authenticity remains verifiable within trust boundaries.

SICF defines identity continuity independent of implementation substrate.

Continuity Outcome

Structural continuity persists despite architectural substitution.

Infrastructure change does not constitute identity termination when invariants are preserved.

D.6.3 Termination and Restoration Attempt

Scenario

A synthetic identity undergoes explicit termination. Later, a prior recorded state is restored and presented as continuation of the terminated identity.

Structural Evaluation

Under Section 9.12 (Lifecycle), termination irreversibly closes active continuity. Restoration following termination constitutes a new synthetic identity with explicit lineage reference to the predecessor.

If restoration is presented as continuation without explicit lineage differentiation, the Non-Silent Lifecycle requirement (8.2.2) fails.

Continuity Outcome

If properly represented as a new identity with lineage reference, continuity semantics remain coherent.

If represented as silent resurrection, continuity degrades into a non-conformant state under Section 9.1.

D.6.4 Assurance Misrepresentation

Scenario

A deployment claims Level 4 continuity assurance without declaring a threat model or specifying adversarial assumptions.

Structural Evaluation

Structural continuity validity may remain intact under Section 9.1.

However, Level 4 assurance requires explicit threat model declaration under Section 12.2.

Without declared threat assumptions, the assurance claim is undefined.

Continuity Outcome

Continuity validity persists.

Continuity strength claim fails.

This distinction reinforces separation between structural invariants and evidentiary robustness.

D.6.5 Expressive Drift and Personality Evolution

Scenario

A synthetic identity operates over extended duration. Its conversational style, tone, and behavioral heuristics evolve significantly. Externally, it appears highly human-like.

Structural Evaluation

Identity continuity remains intact provided:

- Authority anchoring lineage is preserved.
- Lifecycle events remain explicit.
- Continuity remains verifiable.
- Authenticity holds within trust boundaries.
- Synthetic classification remains unambiguous.

Expressive similarity to humans does not redefine identity semantics.

Continuity Outcome

Structural continuity persists independent of personality evolution.

Identity continuity is not equivalent to expressive consistency.

D.6.6 Infrastructure-Bound Identity Reset

Scenario

A synthetic identity is implicitly tied to a vendor-specific credential. During migration, the credential changes and no portable lineage representation exists.

Structural Evaluation

When identity semantics are infrastructure-bound rather than invariant-defined:

- Portability Failure (Section 6.2) emerges.
- Vendor-Bound Identity Collapse (Section 6.5) occurs.
- Non-Silent Lifecycle invariant may fail if transition is indistinguishable from termination and recreation.

Continuity Outcome

Continuity degrades when invariants are not preserved across infrastructure transition.

D.6.7 Institutional Authority Dissolution

Scenario

A principal anchoring a synthetic identity ceases to exist due to bankruptcy, merger, governmental dissolution, or legal restructuring. The synthetic identity remains operational and retains recorded authority anchoring lineage.

Structural Evaluation

Structural continuity remains intact if:

- Authority anchoring lineage remains reconstructable.
- Lifecycle events remain explicit and non-silent.
- Identity authenticity remains verifiable within relevant trust boundaries.

- Synthetic classification remains unambiguous.

SICF defines authority anchoring as a traceable lineage construct. It does not require the continued existence of the originating principal for identity continuity to remain structurally coherent.

Governance-layer frameworks may define authority reassignment, succession, or revocation procedures. These actions do not alter Layer-0 identity continuity invariants.

Continuity Outcome

Identity continuity persists structurally.

Authority legitimacy and succession are governance-layer concerns.

D.7 Structural Sufficiency Under Illustrative Validation

The scenarios above illustrate that the invariants defined in Section 9.1 are sufficient, within the scope and assumptions defined in Section 7, to preserve identity continuity under conditions of:

- Multiplicity
- Migration
- Termination and restoration
- Assurance-level misrepresentation
- Expressive evolution
- Infrastructure transition

Where invariants are preserved, continuity persists.

Where invariants fail, continuity collapses into one or more of the failure classes defined in Section 6.

No additional structural requirements are introduced by these scenarios.