Linked.Art and Related Standards for Time-Based Media Descriptive Metadata

Report produced by Rae Egan, Ontology & Data Modelling Specialist
for **TRANSFER Data Trust**

January 2026

Table of Contents

Background

TRANSFER Data Trust is a decentralized artist-owned archive. The mission is to cooperatively maintain artworks in perpetuity, ensuring their preservation and access across generations. The Trust manages and grows the value of digital assets by leveraging decentralized storage and encryption, backed by a network of care. Initiated in 2023 with early support from Gray Area, the Knight Art + Tech Expansion Fund, Filecoin Foundation for the Decentralized Web (FFDW) and Filecoin Foundation, this innovative project offers a new approach to media art conservation, valuation, and governance.

Cross-Disciplinary Collaboration

In 2024 TRANSFER convened a team of researchers, technologists, conservators, and archivists to support the founding artists in establishing their archives, and develop the technological infrastructure required to sustain the initiative. Founding artists included Huntrezz Janos, Carla Gannis, Lorna Mills, Rosa Menkman, and Eva Papamargariti collaborating on conservation and decentralized tech-enabled archive management workflows with Kelani Nichole, TRANSFER Founder, Wade Wallerstein, TRANSFER Director and Curator at Gray Area, and Regina Harsanyi, Associate Curator of Media Arts at Museum of the Moving Image and Preventive Conservator. The TRANSFER Care Team was composed of Time-based Media Conservation specialists Eddy Colloton, sasha arden, Claudia Röck, Nicholas Kaplan, and Taylor Healy. The infrastructure development was led by Ryan Betts, with support from Andrew Vivash and Mattaniah Aytenfsu.

New Solidarity Exchange Model for Preservation, Ownership, and Value

TRANSFER Data Trust is a proof-of-concept of an open protocol for appraising, preserving, conserving, and exchanging time-based media art. The initiative seeks to establish a cooperative data stewardship model that is co-owned and co-governed by artists. This model allows artists to grow equity in their IP by pooling resources and managed assets, and shifts power to artists by leveraging principles of data sovereignty backed by smart contracts and decentralized storage. Driven by the real needs of contemporary artists, this infrastructure allows them to maintain ownership and agency over their data. The Trust stewards data in a resilient decentralized storage network running on nodes located within the artists' studios, that store encrypted copies of each other's archives for redundancy and safekeeping.

Safeguarding Data in the Wake of AI

At a moment where the proliferation of AI – and the corporate use of personal data to train AI models – is rampant, it's more important than ever for artists and creatives to retain ownership over their data. In traditional art markets, artists often hand over the management of their work and IP to galleries, which sell it to collectors and split the profit, leaving artists with minimal control over their work and its value. With the rise of AI and simulation technologies, it has become increasingly difficult to manage rights and access, and guarantee authenticity of artistic intellectual property. TRANSFER Data Trust changes this dynamic by giving artists agency over their data, and thus their artwork's value, in a decentralized system for data stewardship.

The foundation of this vision relies on structured semantic metadata that enables both human stewardship and machine understanding of digital archives. As artists and institutions increasingly explore the potential of AI models, robust linked data frameworks become essential to allowing these tools to understand, navigate, and learn from archives, while respecting artist ownership and control. This report details how the Linked Art Data Model can provide the semantic foundation necessary to support conservation workflows, underpinning the project's future vision of sovereign AI tooling within decentralized infrastructure.

1. Executive Summary

This report examines the alignment between the Linked Art Data Model and the descriptive metadata requirements for documenting and preserving time-based media (TBM) artworks within the TRANSFER Data Trust's Decentralized Archive. The analysis focuses specifically on descriptive metadata capabilities, establishing the foundation for a broader metadata framework that will ultimately integrate conservation and technical metadata standards in subsequent phases.

The TRANSFER Data Trust is developing a decentralized, artist-owned archival ecosystem built on the InterPlanetary File System (IPFS) to support long-term preservation of time-based media artworks. The Trust Client, a core component of this infrastructure, requires a robust semantic data model to structure metadata for conservation workflows, component tracking, and iterative documentation across diverse artistic practices. This report assesses the extent to which the Linked Art Data Model meets the descriptive metadata requirements for TBM conservation workflows and identifies areas where extensions or complementary approaches are necessary.

While the broader Trust Client metadata framework will require integration of preservation metadata and technical metadata, this analysis establishes the descriptive metadata foundation first.

The analysis was conducted through collaborative discovery sessions with TRANSFER's project lead and conservation team, focusing on current and anticipated workflows for documenting TBM artworks. Requirements were gathered through examination of real-world conservation scenarios and were then mapped against Linked Art's capabilities to identify areas of complete, partial, or absent coverage.

1.1. Key Findings

The Linked Art Data Model provides robust support for core descriptive elements, including artwork identity, creation events, artist attribution, exhibition history, and provenance tracking. Its object-oriented approach and Semantic Web foundation align well with the Trust Client's JSON-LD serialization requirements, supporting interoperability across decentralized environments. However, several gaps emerge when applying Linked Art to TBM documentation needs:

- **Iteration and Version Management:** TBM works often exist in multiple manifestations across different exhibition contexts, software updates, or artistic revisions. Linked Art's current model lacks adequate mechanisms for representing these complex relationships while maintaining work identity.
- **Component-Level Documentation:** TBM works depend on intricate relationships between hardware, software, media files, and consumables that require granular documentation for preservation planning. The model requires extension to capture technical dependencies, track component obsolescence, and address replacement scenarios.
- **Conservation Workflow Integration:** Documentation of conservation decisions, interventions, and ongoing stewardship activities requires more specialized event modeling than Linked Art currently provides for TBM contexts.
- **Multi-Modal Practice Support:** Artists working primarily in TBM often produce related works in installation, sculpture, or hybrid formats that share conceptual and material relationships requiring documentation within the same archival framework.

1.2. Approach

1.2.1. Short-Term Actions

- Develop Linked Art implementation for TBM-specific classes and properties, focusing on iteration modeling and component relationships.
- Create ontology design patterns that can be reused across different TBM conservation scenarios.
- Establish crosswalk mechanisms to preserve interoperability with standard Linked Art implementations.

1.2.2. Long-Term Strategy

- Contribute proposed implementations to the Linked Art community for broader adoption.
- Develop a modular ontology framework that layers TBM-specific capabilities onto Linked Art's foundation.
- Create validation frameworks for testing extension compatibility across decentralized storage environments.

Addressing these gaps is critical for the Trust Client's success in supporting artist-driven preservation practices. The identified extensions will enable comprehensive documentation of TBM works while maintaining semantic interoperability with broader cultural heritage systems. This analysis provides the foundation for developing JSON-LD serialized metadata templates that support real-world conservation workflows within TRANSFER's decentralized infrastructure.

The findings will inform immediate development of Linked Art extension proposals and metadata templates, while establishing a roadmap for integrating preservation and technical metadata standards in future phases. This phased approach ensures thorough coverage of TBM documentation needs while maintaining practical implementation timelines and community engagement opportunities.

2. Context

The Trust Client's metadata framework should rely on Semantic Web technologies to enable interoperability, long-term preservation, and decentralized access to cultural heritage information. This section provides essential context for understanding how descriptive metadata, ontologies, and JSON-LD serialization support the technical requirements of time-based media documentation.

2.1. TRANSFER Trust Client

The Trust Client is a core component of the TRANSFER Data Trust's strategy for building a resilient, artist-owned, and decentralized archive for documenting and preserving art. While its design prioritizes the complex needs of time-based media (TBM) artworks, the Trust Client is equally capable of managing TBM-adjacent works, including sculpture, installation, and other physical or hybrid forms produced by artists whose primary practice is in TBM.

Developed with support from the Knight Foundation and the Filecoin Foundation for the Decentralized Web, the Trust Client is a containerized, peer-to-peer (P2P) file system built on the InterPlanetary File System protocol (IPFS) that enables secure, encrypted-at-rest storage and transfer of digital assets across distributed nodes. Designed to operate independently of commercial cloud infrastructure, the Engine and related technologies ensure long-term accessibility while giving artists and stewards direct control over creative work. For the current phase, content is stored locally across nodes, with a centralized test instance running on Digital Ocean. Future iterations will enable fully decentralized deployments, including private IPFS networks and local hosting for preservation resilience.

A primary function of the Trust Client is to facilitate the creation and management of archival information packages (AIPs); structured sets of digital objects and associated metadata encompassing both descriptive/contextual information and preservation/conservation documentation. While the initial implementation plans to use Linked Art as the core descriptive metadata model, the Trust Client must also address needs that extend beyond basic description, such as:

- Iteration and version tracking across multiple modalities and bodies of intent (e.g., installations, social media deployments, software updates).
- Documentation of components, including hardware, software, consumables, and media files.
- Recording of preservation and conservation events (e.g., migrations, interventions).
- Technical environment details and dependencies.
- Flexible accommodation of TBM and non-TBM works in an artist's practice, ensuring the same archival integrity and interoperability.

Art documentation practices have historically been developed around traditional media such as painting, printmaking, sculpture, and decorative arts, where works typically exist as discrete physical objects with relatively stable material properties. These established frameworks emphasize information about materials, dimensions, condition, and provenance that align well with the characteristics of conventional art forms. However, documenting contemporary works, particularly time-based media, requires fundamentally different information that does not map cleanly to existing schemas and standards developed for traditional media.

Time-based media works challenge conventional documentation approaches through their dependence on technology, their existence across multiple iterations and formats, their temporal nature, and their complex relationships between conceptual intent and technical implementation. These characteristics require metadata frameworks that can capture not only what a work is, but also how it behaves, evolves, and maintains its identity across different manifestation contexts.

2.2. Descriptive Metadata for Art Documentation

Descriptive metadata for art documentation encompasses the essential information needed to identify, discover, and contextualize artworks within cultural and historical frameworks. It typically includes core elements such as artwork titles, artist attribution, creation dates, media and materials, dimensions, exhibition history, and provenance information. This metadata serves multiple constituencies: artists, their studios, and their estates require accurate documentation of creative intent and work identity; conservators need contextual information to inform preservation decisions; researchers depend on descriptive metadata for discovery and analysis; gallerists and collectors use it for valuation, exhibition planning, and market documentation; and institutions use it to manage collections and facilitate public access.

In traditional art documentation, descriptive metadata has often been managed in isolated systems with limited interoperability. The distributed nature of contemporary art practices—particularly time-based media works that may exist across multiple institutions, formats, and exhibition contexts—requires more flexible and connected approaches to metadata management that can accommodate both the complexity of the works themselves and the collaborative nature of their preservation.

This report examines how the Trust Client's descriptive metadata requirements for TBM and TBM-adjacent works align with the capabilities of the Linked Art Data Model. While the broader Trust Client metadata framework may ultimately require integration of preservation metadata (e.g., PREMIS) and technical metadata (e.g., EBUCore, PBCore), this analysis establishes the descriptive metadata foundation first. This phased approach enables thorough examination of each metadata domain before developing an integrated framework that supports the complex, interdisciplinary requirements of TBM preservation.

2.3. Semantic Technology Fundamentals

2.3.1. Ontologies

The term ontology is ambiguous with different meanings in different fields, such as computer science, linguistics, information science, and philosophy. To complicate things further, there is no clear distinction between "ontologies" and "vocabularies" in the context of the Semantic Web (Hyvönen, 2012). For this project, an ontology can be defined as a formal and explicit description of concepts in a domain of discourse (Noy & McGuinness, 2001). To put it more simply, an ontology is a schema or conceptual framework that gives data meaning.

2.3.2. The Semantic Web

In computing, Linked Data describes a method of publishing structured data using standard web technologies, such as Hypertext Transfer Protocol (HTTP), Resource Description Framework (RDF), and Uniform Resource Identifiers (URIs), so that resources can be interlinked for access by both humans and machines. While the World Wide Web (WWW) links associated pages with each other for human use, the Semantic Web, driven by Linked Data, links underlying concepts and machine-processable data sources together.

The universal framework for the Semantic Web is RDF, a standard for modeling data that uses three simple components: a subject, predicate, and object (Taylor, 2015, p. 4). For example, to express "Lynn Hershman Leeson created Agent Ruby," the subject is "Lynn Hershman Lesson," the predicate is "created," and the object is "Agent Ruby." For this reason, people often call RDF data triples. Similarly, a triplestore or RDF store is a "purpose-built database for the storage and retrieval of triples through semantic queries" (Triplestore, n.d.). When triples are linked, they form an abstract data type called a graph structure. RDF has many syntaxes, such as RDF/XML, N-Triples, Terse RDF Triple Language (Turtle), RDFa, and JSON-LD. The choice of these textual syntaxes, called serializations, depends on the use case.

2.3.3. JSON-LD Serialization

JSON-LD (JavaScript Object Notation for Linked Data) bridges the gap between traditional web development practices and Semantic Web technologies. As a serialization format, JSON-LD allows Semantic Web data to be expressed using familiar JSON syntax while maintaining the formal semantics needed for interoperability and automated processing. To illustrate the difference: traditional JSON might represent a work as `{"title": "Agent Ruby", "creator": "Lynn Hershman Lesson"}`, which a computer interprets as simple text values. JSON-LD adds context that tells the computer what "title" and "creator" mean semantically by linking them to formal definitions that enable different systems to understand that "creator" refers to the person who made the artwork, not just any arbitrary relationship.

For the Trust Client, JSON-LD offers several critical advantages. The format integrates seamlessly with modern web applications and APIs, making it accessible to developers who may not have extensive experience with the Semantic Web technologies. JSON-LD documents can be processed both as regular JSON for application development and as semantic data for integration with cultural heritage systems. This dual compatibility is essential for the Trust Client's goal of supporting both artist-centered workflows and institutional interoperability.

JSON-LD's context mechanism allows metadata to reference established vocabularies and ontologies while maintaining local flexibility. This means that Trust Client metadata can align with standards like Linked Art while accommodating time-based media-specific extensions without breaking compatibility with existing systems.

2.3.4. Integration with Decentralized Infrastructure

The combination of ontologies and JSON-LD serialization aligns particularly well with the Trust Client's decentralized architecture. Unlike proprietary metadata formats that require specific software or centralized services, JSON-LD documents can be stored, transmitted, and processed across distributed IPFS networks without loss of semantic meaning. This approach ensures that metadata remains accessible and interpretable even as individual nodes join or leave the network.

The semantic structure provided by ontologies also supports resilient metadata management in decentralized environments. When metadata follows formal ontological patterns, automated systems can validate consistency, detect missing information, and maintain relationships between distributed resources (i.e., capabilities that are essential for long-term preservation in peer-to-peer networks). This technical foundation enables the Trust Client to support sophisticated time-based media documentation requirements while maintaining the flexibility and autonomy that are central to TRANSFER's artist-driven approach to digital preservation.

2.3.5. Practical Benefits for Time-Based Media Documentation

For the Trust Client's time-based media documentation needs, semantic technologies provide three essential capabilities that traditional database approaches cannot easily achieve:

- **Flexibility:** Ontologies can accommodate the diverse and evolving nature of TBM artworks without requiring rigid database schemas. When artists develop new forms of practice or conservation needs change, the ontological framework can be extended without breaking existing documentation.
- **Interoperability:** By using shared vocabularies and formal semantic structures, metadata created by individual artists or studios can be understood and integrated by institutions, conservators, and researchers using different systems—supporting the collaborative nature of TBM preservation.
- **Long-term preservation:** Unlike proprietary formats that depend on specific software, semantic data expressed in standards-based formats like JSON-LD can be stored, transmitted, and interpreted across distributed networks for decades without loss of meaning, aligning with the Trust Client's preservation mission.

These capabilities make semantic technologies particularly well-suited for documenting artworks that exist across multiple manifestations, depend on complex technical infrastructures, and require sustained collaborative stewardship over time.

2.4. Semantic Standards for Art Documentation

The cultural heritage community has developed several semantic standards and vocabularies specifically for art documentation that provide structured approaches to describing cultural objects and their relationships. Understanding these standards is essential for positioning the Trust Client's metadata framework within the broader ecosystem of cultural heritage technology.

2.4.1. Comité International pour la Documentation Conceptual Reference Model (CIDOC-CRM)

The CIDOC Conceptual Reference Model serves as the foundational ontology for cultural heritage information. Developed by the International Council of Museums (ICOM) Committee for Documentation, CIDOC-CRM provides a formal structure for describing the implicit and explicit concepts and relationships used in cultural heritage documentation. The model focuses on events as the primary mechanism for connecting people, places, objects, and concepts over time, making it particularly powerful for documenting complex provenance and contextual relationships.

2.4.2. Linked Art Data Model

Linked Art emerged from practical challenges encountered during early adoption of CIDOC-CRM for digital applications. Between 2014 and 2017, three significant U.S.-based projects, Art Tracks at Carnegie Museum of Art, the Getty Provenance Index Remodel, and the American Art Collaborative (AAC), began developing

RDF-native applications that would not only publish but actively consume CIDOC-CRM data. These projects revealed two critical requirements. First, without documented, consistent patterns for applying CIDOC-CRM, different interpretations would produce structurally incompatible expressions of similar information. Second, semantically correct modeling often lacked sufficient precision for "roundtripping" information, making it impossible for software to distinguish between structurally identical but semantically different constructs (Newbury, 2018).

The publication of JSON-LD serialization for RDF in 2014 provided a crucial technical foundation, enabling RDF concepts to be expressed using familiar JSON syntax without requiring developers to understand RDF formalisms. The success of the International Image Interoperability Framework (IIIF) demonstrated that a JSON-LD standard could achieve broad acceptance in cultural heritage by using JSON-LD to hide RDF complexity from implementers while preserving underlying semantic structure (Newbury, 2018).

Linked Art addresses these challenges by using a streamlined profile of CIDOC-CRM as its core ontology, incorporating Getty Vocabularies as primary sources of identity, and JSON-LD as the target serialization format. This approach balances semantic rigor with practical implementation needs, making it more accessible to developers while preserving the conceptual sophistication needed for complex cultural objects.

Key characteristics of the Linked Art Data Model include:

- **Event-driven approach:** Following CIDOC-CRM's emphasis on events to connect entities and track changes over time.
- **JSON-LD native:** Designed specifically for web-based implementation with clear serialization patterns and developer-friendly syntax.
- **Practical patterns:** Provides concrete examples and templates for common art documentation scenarios such as artwork creation, exhibition, ownership, and physical characteristics.
- **Extensible framework:** Supports domain-specific extensions while maintaining core interoperability with standard CIDOC-CRM implementations.
- **Reduced complexity:** Covers the majority of institutional use cases with a fraction of the full CRM's conceptual overhead.
- **Community-driven development:** Actively maintained by a community of cultural heritage professionals and technologists.

2.4.3. Getty Vocabularies

Getty Vocabularies, including the Art & Architecture Thesaurus (AAT), Union List of Artist Names (ULAN), and Getty Thesaurus of Geographic Names (TGN), provide controlled terminology and authority records that support consistent identification and description across cultural heritage systems. These vocabularies integrate seamlessly with semantic frameworks like CIDOC-CRM and Linked Art, providing stable URIs and hierarchical concept structures that enable precise and interoperable metadata creation.

2.4.4. PREservation Metadata: Implementation Strategies (PREMIS)

PREMIS was developed by the Library of Congress as a data model for preservation metadata in digital preservation systems (PREMIS Editorial Committee, 2015). While PREMIS is not traditionally considered a descriptive metadata standard and is rarely applied to art documentation contexts, its robust framework for modeling complex digital objects and their technical relationships offers significant value for time-based media preservation workflows. Its core entity model includes Intellectual Entity (the conceptual work), Representation (specific instances or versions), File (technical manifestations), and Bitstream (sequences of bits within files). This layered approach to modeling digital content provides granular control over technical documentation while maintaining connections to conceptual identity.

2.4.5. Standards Limitations for Time-Based Media

Together, these standards create a robust foundation for semantic art documentation that supports both machine processing and human understanding. However, their separate development trajectories reveal complementary gaps when applied to contemporary art forms like time-based media that challenge conventional documentation approaches.

CIDOC-CRM, Linked Art, and Getty Vocabularies were primarily developed to support museums, galleries, and cultural heritage institutions managing accessioned collections. This institutional focus shapes their priorities around cataloging, public access, and retrospective documentation of works that have entered permanent collections. While their event-based modeling and conceptual rigor provide strong foundations for documenting provenance, creation, and exhibition contexts, they lack native support for the technical complexity, iterative development, and hybrid materiality that characterize TBM works.

The Trust Client, by contrast, serves a fundamentally different documentation context (e.g., living artists, their studios and archives, active conservation practices, art market considerations). This context requires documentation frameworks that support ongoing artistic development, edition management, valuation and distribution, artist-driven preservation planning, and studio workflow integration. Artists and their teams need to document and track technical iterations across exhibition contexts, manage component inventories and maintenance schedules, and maintain detailed records that extend beyond the retrospective cataloging focus of institutional collection management systems.

Linked Art's streamlined approach to CIDOC-CRM improves accessibility but inherits both the institutional collection management focus and technical limitations around versioning, component dependencies, and the strict physical/digital object distinctions that complicate documentation of works existing across multiple modalities simultaneously. The model's emphasis on "90 percent of the use cases of 90 percent of the organizations" (Linked Art Editorial Board, n.d., *Profile*). reflects institutional collection contexts rather than artist-studio or art market scenarios where edition structures, ongoing technical development, and preservation planning occur prior to, and independently of, institutional acquisition. Similarly, Getty Vocabularies offer essential controlled terminology but require extension with TBM-specific concepts for software types, technical processes, digital components, and contemporary artistic practices that have emerged since these vocabularies' initial development.

These complementary limitations reveal that comprehensive TBM documentation requires integration across standards rather than reliance on any single framework. The technical preservation capabilities of PREMIS can address gaps in Linked Art's component and versioning support, while Linked Art's art-historical sophistication can provide the contextual documentation absent from PREMIS. However, this integration itself presents challenges around conceptual alignment, compatibility, and maintaining semantic coherence across multiple standards designed for different domains.

3. Model Development Methodology

This report was developed through a collaborative and iterative discovery and evaluation process. Biweekly sessions with the project lead and members of the care team served as the primary forum for gathering information. These sessions focused on documenting current and anticipated conservation workflows, reviewing example works from TRANSFER, and discussing practical considerations for appraisal, preservation, and documentation.

3.1. Workflow Selection and Iterative Model Development

The Trust Client supports multiple conservation workflows for time-based media documentation. To focus this initial report, the iteration reporting workflow was prioritized. Iteration reports provide descriptive metadata for specific installations, helping track the evolution of artworks over time, and providing parameters that can be used to update other conservation workflows.

The conceptual data model supporting the iteration reporting workflow was developed through three cycles of iterative design and review with stakeholders to define core classes, relationships, and cardinality constraints. Each cycle involved eliciting functional requirements through the examination of specific TBM artworks and iteration reporting scenarios, as well as model refinement based on stakeholder feedback and real-world workflow validation. This approach ensured that the model remained grounded in practical conservation needs while identifying areas where standard descriptive metadata frameworks require extension for TBM documentation.

3.2. Class-Level Focus and Modular Design

This report focuses exclusively on modeling and mapping classes (i.e., entities) rather than properties (i.e., relationships). In ontological terms, classes represent categories or types of things that exist in the domain, such as artworks, people, or events, while properties define the relationships and attributes that connect or describe these entities (Noy & McGuinness, 2001).

While the analysis defines basic relationships and cardinality constraints between classes (e.g., "one Work can have multiple Editions"), detailed descriptive properties such as title, medium, duration, or dimensions are not modeled. This level of property modeling represents a more granular analysis that is better addressed once class-level coverage gaps are identified and resolved. Where suitable class matches in Linked Art are identified, TRANSFER can leverage their existing properties and extend with additional properties as needed for TBM-specific requirements.

This approach follows established ontology engineering methodologies on multiple levels. First, it implements a top-down development strategy, which starts by defining the most general concepts in the domain before specializing them further (Uschold & Gruninger, 1996). Second, it aligns with modular ontology design principles, where individual classes correspond to discrete ontology modules that can be developed, maintained, and reused independently (Hitzler et al., 2017). In this framework, each core class identified in our analysis, such as Work, Component, or Iteration, represents a potential ontology module that encapsulates the concepts and relationships specific to that entity type.

This modular, class-focused methodology provides several analytical advantages. It enables assessment of coverage gaps at the module level, ensuring that fundamental entity types can be accommodated before addressing the more complex question of detailed property modeling. It supports efficient identification of structural gaps in existing standards while providing a foundation for more detailed property-level analysis in subsequent phases. Most importantly, it establishes whether Linked Art can accommodate the core entity types required for iteration reporting, allowing us to determine the scope of extensions needed before addressing granular descriptive properties.

The methodology for mapping these classes to external standards is detailed in Section 5.1.

4. TRANSFER Data Trust Conceptual Data Model

4.1. Core Class Definitions

Activity

An Activity is an event or process through which something is created, configured, exhibited, installed, migrated, or maintained. Iterations are modeled as specific types of Activities. Activities are carried out by one or more Agents in specific roles.

Agent

An Agent is a person, group, or organization responsible for carrying out an Activity. Agents can play various roles, including, but not limited to: artist, conservator, curator, gallerist, developer, and fabricator. Roles can be qualified and recorded in relation to the Activities they participate in.

Component

A Component is a discrete, configurable, or reusable part of a Variant required by a Version or Iteration. Components may be software, hardware elements, display surfaces, media assets, scripts, or physical parts required for execution, exhibition, or preservation of a Variant. Components are not Works on their own, but may be shared across multiple Variants, Versions, or Iterations. Components may be shared, replaced, updated, or maintained independently and are often the smallest documented unit in a conservation workflow.

Edition

An Edition is a unit of a Work with an appraisal value. It may specify rights, quantity, or terms of use, is often limited or numbered, and may group one or more identical or similar Variants together. Editions are the primary organizing unit for distribution and ownership. One Work may have one or multiple Editions. Editions can be unique (i.e., 1/1), part of a limited run, or open, and can include artist's proofs.

Iteration

An Iteration is a time- and site-specific execution, installation, or exhibition of a Variant. It is modeled as a subclass of Activity, and captures when, where, and how the Variant was instantiated. Iterations can reference one or more Versions and use multiple Components. One Variant may have multiple Iterations.

Series

A Series is a conceptual grouping of related artworks, typically conceived by an artist. Works within a Series may share thematic, formal, or procedural continuity but are still distinct artworks.

Variant

A Variant defines the form factor or modality of a Work within an Edition. It describes how the artwork is configured or intended to be experienced (e.g., projection, sculptural, single-channel video, web-based). A Variant groups together Versions and Iterations that share the same modality. One Edition can have multiple Variants.

Version

A Version represents a technical or procedural update to the Variant, such as software changes, hardware substitutions, or encoding differences. Versions are created serially over time, documenting how the artwork has been altered without changing its conceptual identity. One Variant may have multiple Versions.

Work

A Work is an abstract conception of an artwork, encompassing at minimum its artistic intent and identity across different expressions, configurations, and realizations. It is the intellectual content that persists across Editions and Variants. It may be part of a Series (or none) and may have zero, one, or more Editions.

4.2. Class Relationships and Cardinality

Source Entity Label	Relationship Label	Target Entity Label	Relationship Cardinality	Relationship Commentary
Series	references	Series	M:M	Captures thematic or procedural connections between different Series, documenting artistic lineages, influences, or ongoing conceptual dialogues. A Series may reference zero, one, or many other Series, and may be referenced by multiple other Series.
Series	includes	Work	0:M	Establishes membership of Works within a Series. A Series may include zero, one, or many Works, but a Work belongs to at most one Series (or none).
Work	references	Work	M:M	Documents intellectual or aesthetic connections through citation, appropriation, or conceptual dialogue. A Work may reference zero, one, or many other Works, and may be referenced by multiple Works.
Work	hasEdition	Edition	0:M	Instantiates a Work for appraisal, distribution, ownership, etc. A Work may have zero, one, or many Editions, with each Edition belonging to exactly one Work. Editions may have different rights, quantities, or terms of use.
Edition	hasVariant	Variant	0:M	Defines different modalities or form factors within the same Edition. An Edition may have zero, one, or many Variants (e.g., projection vs. sculptural vs. web-based), with each Variant belonging to exactly one Edition.
Variant	hasVersion	Version	0:M	Documents technical or procedural updates without changing fundamental modality. A Variant may have zero, one, or many Versions, with each Version belonging to exactly one Variant. Versions are created serially over time as technical changes occur.
Version	derivesFrom	Version	M:M	Captures technical evolution and serial development over time. A Version may derive from zero, one, or many previous Versions and may serve as the basis for multiple subsequent Versions in the serial progression.
Variant	hasIteration	Iteration	M:M	Instantiates a Variant across different times and places (i.e., installation). A Variant may have zero, one, or many Iterations, and an Iteration may instantiate one or many Variants simultaneously in complex installations.
Iteration	iteratesOn	Iteration	M:M	Documents how installations build upon previous exhibition experiences. An Iteration may iterate on zero, one, or many previous Iterations, and may serve as the basis for multiple future Iterations, capturing learning and refinement across contexts.
Variant	requires	Component	M:M	Specifies technical dependencies for realization. A Variant requires one or many Components, and a Component may be required by multiple Variants, enabling documentation of component reuse.
Component	dependsOn	Component	M:M	Documents technical dependencies crucial for preservation planning. A Component may depend on zero, one, or many other Components (e.g., software requiring specific hardware), and may be depended upon by multiple other Components.

Version	usedIn	Iteration	M:M	Tracks which technical configurations have been deployed. A Version may be used in zero, one, or many Iterations, and an Iteration may employ one or many Versions when different aspects use different technical implementations.
Version	specifies	Component	M:M	Documents technical requirements at the Version level. A Version may specify zero, one, or many Components as required or recommended, and a Component may be specified by multiple Versions.
Iteration	uses	Component	M:M	Captures Components used in installations. An Iteration uses one or many Components, and a Component may be used by multiple Iterations. Actual usage may differ from requirements due to availability or substitution.
Activity	conceives	Series	1:M	An Activity may conceive zero, one, or many Series, but each Series is conceived through exactly one Activity. Future state example: ActivityType (ConceptionActivity), AgentRole (creator).
Activity	creates	Work	1:M	Establishes creative authorship. An Activity may create zero, one, or many Works, but each Work is created through exactly one creation Activity. Future state example: ActivityType (CreationActivity), AgentRole (creator).
Activity	produces	Edition	1:M	Documents edition processes. An Activity may produce zero, one, or many Editions, but each Edition is produced through exactly one production Activity. Future state example: ActivityType (ProductionActivity), AgentRole (producer, fabricator).
Activity	defines	Variant	M:M	Captures specification activities. An Activity may define zero, one, or many Variants, and a Variant may be defined through one or many Activities involving different stakeholders (e.g., artist specifications, curatorial decisions, technical requirements). Future state example: ActivityType (CurationActivity, TechnicalActivity), AgentRole (curator, technical-specialist).
Activity	versions	Version	M:M	Documents versioning processes that occur serially over time. An Activity may create zero, one, or many Versions, and a Version may result from one or many Activities (e.g., testing, implementation, documentation within the serial progression). Future state example: ActivityType (TechnicalActivity), AgentRole (conservator, developer, tester).
Activity	isA	Iteration	1:1 (inheritance)	Establishes that Iteration is a specialized type of Activity. Each Iteration is exactly one Activity, and this Activity is exactly one Iteration, with additional temporal and spatial specificity. Future state example: ActivityType (ExhibitionActivity), AgentRole (installer, curator).
Activity	maintains	Component	M:M	Documents maintenance processes essential for conservation tracking at the Component level. An Activity may maintain zero, one, or many Components (e.g., system-wide updates), and a Component may be maintained through multiple Activities over time. Future state example: ActivityType (ConservationActivity), AgentRole (conservator, technician).

Agent	carriesOut	Activity	M:M	Establishes responsibility and tracks expertise across workflows. An Agent may carry out zero, one, or many Activities in various roles, and an Activity may involve one or many Agents (e.g., collaborative creation, conservation teams). Should be refined to include AgentRole intermediate class, ActivityType classification, and RoleType controlled vocabulary (creator, collaborator, curator, conservator, installer, etc.) in future state.

4.3. Mermaid Diagram

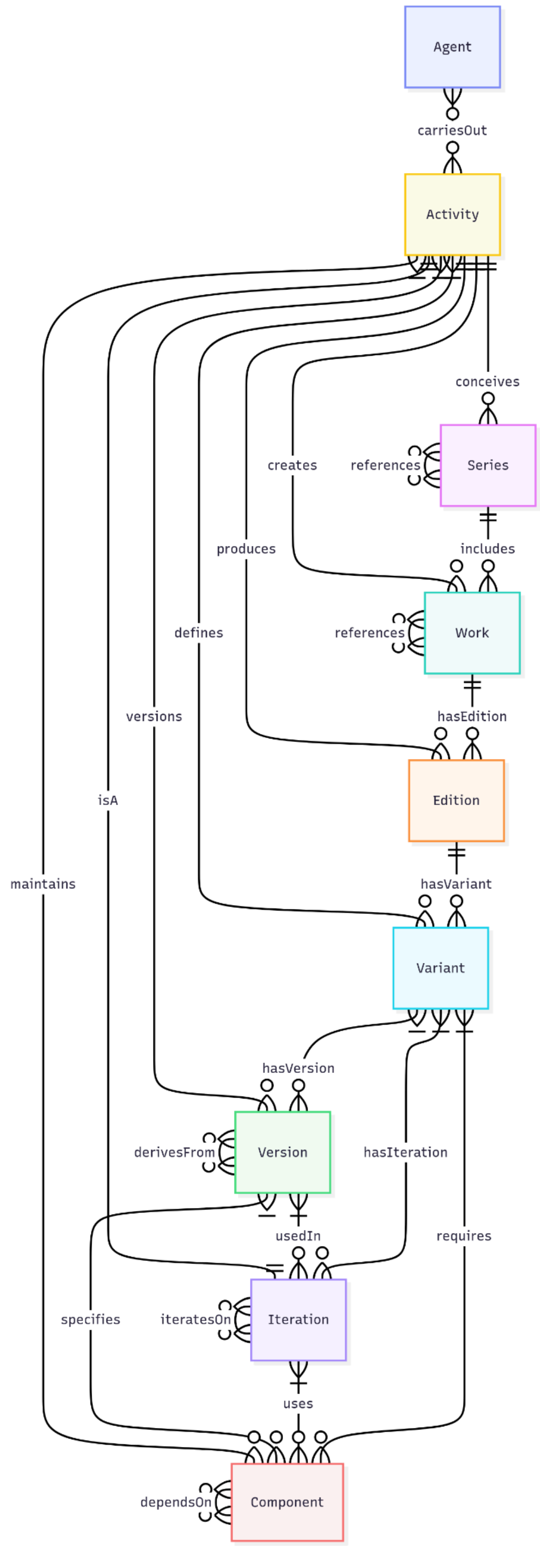

4.4. Future State Considerations

The current `Activity` and `Agent` modeling represents a foundational approach that establishes core structural relationships while acknowledging that more sophisticated role-based modeling will be developed in subsequent phases. This simplified approach enables completion of the class-level report while providing a solid foundation for future refinement.

In future modeling phases, the `Activity` class should be expanded to include controlled vocabulary for activity types (e.g., CreationActivity, CurationActivity, ConservationActivity, ExhibitionActivity, TechnicalActivity). Similarly, the direct `Agent`-`Activity` relationship will be refined through an intermediate AgentRole class that captures specific role assignments (e.g., creator, collaborator, curator, conservator, installer, fabricator).

5. Analysis

5.1. Semantic Mapping Methodology and Principles

5.1.1. Rationale for Standards Interlinking

Semantic interoperability depends on establishing explicit relationships between classes in heterogeneous semantic models. By creating formal mappings between TRANSFER's conceptual model and established standards like Linked Art and PREMIS, the Trust Client can maintain compatibility with broader cultural heritage ecosystems while accommodating time-based media-specific requirements. Recommended primary mappings can be used to link classes that belong in heterogeneous semantic models. Such links enable semantic interoperability between applications (Alexopoulos, 2020, p. 28). The following benefits of interlinking outweigh potential maintenance costs:

- **Institutional Interoperability:** CIDOC-CRM and its Linked Art profile are widely adopted across museums, galleries, and cultural heritage institutions. Formal mappings enable TRANSFER metadata to be understood and integrated by institutional systems without requiring custom translation layers.
- **Preservation Community Alignment:** PREMIS serves as the standard for preservation metadata in many digital repositories worldwide. Mapping to PREMIS ensures that Trust Client documentation aligns with established preservation practices and can be exchanged with other preservation systems.
- **Knowledge Reuse:** Rather than developing documentation practices in isolation, formal mappings allow TRANSFER to leverage existing vocabularies, best practices, and implementation guidance from mature standards communities.
- **Future-Proofing:** As Linked Art and PREMIS evolve, maintaining explicit mappings provides a framework for assessing how changes in external standards affect the Trust Client's data model and implementation.

5.1.2. Mapping Principles

To ensure semantic accuracy and maintainable interlinking, the following principles guide all class-level mapping decisions.

Principal 1: Minimum Generalization

Mappings should be as specific as possible without introducing semantic ambiguity. Overly broad mappings sacrifice precision and limit the practical value of interoperability. Conversely, forcing mappings to inappropriately narrow classes introduces semantic drift that can mislead downstream consumers of the metadata.

Principal 2: Single Primary Mapping

Each TRANSFER class should have at most one primary mapping to external standards to reduce maintenance burden and prevent conflicting semantic assertions. When both Linked Art and PREMIS offer equivalent classes, always prioritize Linked Art. This preference for Linked Art reflects several critical considerations:

- **Domain Alignment:** Linked Art was specifically designed for art documentation, making its conceptual framework naturally aligned with TRANSFER's requirements for capturing artistic intent, creative processes, exhibition contexts, and aesthetic relationships. PREMIS, while excellent for technical preservation, was designed for library and archival contexts and requires significantly more semantic assertions to adequately capture artistic contexts.
- **Semantic Coherence:** Maintaining mappings to multiple distinct semantic models creates substantial complexity in implementation, validation, and long-term maintenance. Prioritizing a single primary standard (i.e., Linked Art) wherever possible reduces this complexity and creates a more coherent metadata framework.
- **Community Alignment:** The Linked Art community includes museums, galleries, and contemporary art institutions who may face similar TBM documentation challenges. Aligning primarily with Linked Art facilitates knowledge sharing, community contribution, and potential adoption of TRANSFER extensions by other institutions.

Principal 3: Exact Class-Level Matches Only

Following the methodology established in Section 3.2, this analysis focuses exclusively on class-level mappings, asserting relationships only when classes are semantically equivalent or when a clear subsumption relationship exists. For example, we avoid using skos:closeMatch or similar approximate alignments.

Property-level mappings should be addressed in subsequent phases after class-level gaps are resolved and TRANSFER-specific extensions are developed. Where suitable class matches are identified in Linked Art, TRANSFER can leverage their existing properties and extend with additional properties as needed for TBM-specific requirements.

Exception: Type-Based Disambiguation in Linked Art

The mapping principles above assume relatively symmetric class structures between source and target ontologies. However, a fundamental architectural distinction between the TRANSFER conceptual model and Linked Art qualifies how these principles apply in practice.

The TRANSFER model employs explicit, discrete classes for each core entity where class membership alone conveys semantic meaning. Linked Art, by contrast, reuses broad base classes from CIDOC CRM, distinguishing many of them through rdf:type and crm:P2_has_type. This type-based disambiguation pattern means that multiple conceptually distinct TRANSFER classes may map to the same Linked Art base class, differentiated only through classification assertions. For example, both Series and Edition could map to la:Set but would need to be distinguished through classification with distinct AAT terms (e.g., aat:300027349 for Series, aat:300121294 for Edition). As a result, while this analysis focuses on class-level mappings per the methodology established in Section 3.2, some mappings presented in Section 5.2 are incomplete without their accompanying classification patterns documented in the property-level implementation, and have been included in commentary.

5.2. Class-Level Mapping

TRANSFER Core Class	Linked Art Mapping	CIDOC CRM Mapping	Linked Art/CIDOC CRM Coverage	Linked Art/CIDOC CRM Coverage Commentary	PREMIS Mapping	PREMIS Coverage	PREMIS Coverage Commentary	Recommendations
Series	Set	N/A	Partial Coverage	Linked Art's Set class provides structural support for groupings, with type disambiguation achieved through classification with AAT term aat:300027349 (Series) under aat:300241507 (Object Grouping)	Intellectual Entity	Partial Coverage	PREMIS Intellectual Entity can represent conceptual groupings, but lacks explicit semantics for artist-conceived series, thematic continuity, or procedural relationships. Custom properties may be needed to capture these series-specific semantics, depending on stakeholder requirements for documenting artistic intent.	- Map Series to Set with classified_as aat:300027349 (Series) - Consider proposing new Series class to Linked Art; if approved and implemented, map Series to Series.
Work	PropositionalObject	E89 Propositional Object	Partial Coverage	LinkedArt PropositionalObject maps to a single conceptual work identity, with some hierarchical constraints due to CIDOC-CRM E71 Human-Made Thing.	Intellectual Entity	Strong Coverage	PREMIS Intellectual Entity provides alternative mapping for conceptual-level identity without CIDOC-CRM E71 Human-Made Thing constraints.	- Map Work to PropositionalObject
Edition	Set	No direct equivalent	Partial Coverage	Linked Art's Set class provides structural support for edition groupings, with type disambiguation achieved through classification with AAT term aat:300121294 (Edition) under aat:300241507 (Object Grouping)	No direct equivalent	Gap	Lacks concepts for appraisal value, ownership units, edition numbering, artist's proofs, etc.	- Map Edition to Set with classified_as aat:300121294 (Edition) - Consider proposing new Edition class to Linked Art; once approved and implemented, map Edition to Edition
Variant	HumanMadeObject	E22 Human-Made Object	Partial Coverage	Linked Art HumanMadeObject maps to a single form factor with type disambiguation through local terminology	Representation	Strong Coverage	PREMIS Representation supports hybrid materiality as "a single instance of an Intellectual Entity held in a preservation repository."	- Map Variant to HumanMadeObject or DigitalObject with classified_as using local term "Variant" - Consider proposing new

								Variant class to Linked Art; once approved and implemented, map Edition to Edition
Version	InformationObject	E73 Information Object	Partial Coverage	Linked Art InformationObject maps to Version with type disambiguation through classification using AAT term aat:300121393 (Version)	agentVersion and environmentalVersion (no direct class equivalent)	Partial Coverage	PREMIS provides agentVersion (for software/hardware updates) and environmentalVersion (for changes in technical environment).	- Map Version to InformationObject with classified_as aat:300121393 (Version) - Consider proposing new Version class to Linked Art; if approved and implemented, map Version to Version
Iteration	Activity > Exhibition	E7 Activity	Partial Coverage	Strong mapping to Linked Art's Activity > Exhibition but with a potentially large semantic gap between "exhibition" and the established concept of an "iteration" in a TBM conservation sense; requires stakeholder decision on terminology alignment	Event	Partial Coverage	PREMIS Event captures preservation actions but lacks exhibition/installation semantics and spatial/temporal context specific to iteration execution	- Map Iteration to Activity
Component	HumanMadeObject / DigitalObject	E22 Human-Made Object / D1 Digital Object	Partial Coverage	Linked Art HumanMadeObject or DigitalObject maps to Component materiality, with type disambiguation through classification using AAT term aat:300241490 (Component)	File / Object / Environment	Partial Coverage	PREMIS File and Object can represent digital components (media assets, software), and Environment entities can describe hardware infrastructure. However, PREMIS lacks semantics for sculptural objects, fabricated display surfaces, and physical components with aesthetic significance. It does not capture material properties, dimensions, spatial relationships, or the artistic/conservation importance of physical components beyond their technical	- Map Component to HumanMadeObject or DigitalObject with classified_as aat:300241490 (Component) - Document guidelines for selecting HumanMadeObject vs. DigitalObject based on materiality - Consider proposing new Component class to Linked Art; if approved and

							function.	implemented, map Component to Component - Consider creating custom type property (e.g., componentType) using controlled vocabulary to describe materiality and function
Activity	Activity	E7 Activity	Strong Coverage	Direct alignment with Linked Art Activity and CIDOC-CRM E7 Activity; supports diverse activity types including creation, modification, exhibition, and conservation.	Event	Strong Coverage	PREMIS Event provides strong alignment for preservation-related activities, with controlled vocabulary for event types (e.g., migration, fixity check, validation).	- Map Activity to Activity - Consider developing controlled vocabulary for TRANSFER-specific Activity types
Agent	Actor (Person / Group)	E39 Actor	Strong Coverage	Full alignment with Linked Art Actor (e.g., Person, Group) and CIDOC-CRM E39 Actor; supports role attribution through participation in activities.	Agent	Strong Coverage	PREMIS Agent aligns well with organizational and individual agents responsible for preservation events, supporting agent types (e.g., person, organization, software).	- Map Agent to Actor

5.3. Coverage Analysis

5.3.1. Strong Coverage Areas

Work

The mapping of TRANSFER's Work class to Linked Art's PropositionalObject provides a conceptual-level work identity that captures artistic intent before physical or technical manifestation. However, this mapping requires consideration of the underlying CIDOC-CRM hierarchy. E89 Propositional Object is a subclass of E28 Conceptual Object, which is itself a subclass of E71 Human-Made Thing, defined as "discrete, identifiable human-made items that are documented as single units" (CIDOC CRM Special Interest Group, 2024).

This hierarchical relationship means that using PropositionalObject for TRANSFER's Work class assumes that all TBM works qualify as human-made things. While this assumption holds for most contemporary art practices, there may be potential future complexity in TBM contexts where works may incorporate algorithmic processes that generate content autonomously, machine learning systems that produce outputs beyond direct human control, networked interactions where emergent behaviors arise from system dynamics rather than explicit artistic programming, or AI-generated components that may not meet traditional definitions of human-made items. For the present, TRANSFER stakeholders are comfortable asserting that all Works can be classified as human-made, recognizing that this classification may require revisiting as artistic practices evolve.

PREMIS offers an alternative through its Intellectual Entity concept, which represents works at the conceptual level without the ontological constraints of E71. This provides greater flexibility for documenting works that may involve computational generation, artificial intelligence, or other non-traditional creative processes. However, following the Mapping Principles established in Section 5.1.2, which prioritize Linked Art for domain alignment and semantic coherence, TRANSFER should use Linked Art's PropositionalObject as the primary mapping despite any theoretical advantages of PREMIS Intellectual Entity's greater flexibility.

Recommendations

- Map TRANSFER Work to Linked Art PropositionalEntity via rowl:equivalentClass as the primary mapping.

Activity

Both Linked Art and PREMIS provide strong native support for modeling activities and events. Linked Art's Activity class maps directly to CIDOC-CRM E7 Activity, supporting diverse activity types including creation, modification, exhibition, and conservation. PREMIS's Event class offers parallel functionality with an emphasis on preservation-related actions and controlled vocabulary for event types such as migration, fixity checking, and validation. This strong coverage provides flexibility for different aspects of TBM documentation, with Linked Art for art-related activities (e.g., creation, exhibition) and PREMIS for technical preservation activities (e.g., migration, validation).

Recommendations

- Map TRANSFER Activity to Linked Art Activity as the primary mapping.
- Develop controlled vocabulary for TRANSFER-specific Activity types that extends both Linked Art and PREMIS event type vocabularies and establish guidelines on usage.
- Consider future refinement to include activity type classification and role-based modeling as noted in Section 4.4 Future State Considerations.

Agent

Both frameworks provide comprehensive support for documenting people and organizations responsible for activities. Linked Art's Actor class, encompassing Person and Group, aligns directly with CIDOC-CRM E39 Actor, while PREMIS's Agent class supports human, machine, and software types. Role attribution is well-supported through participation in activities in both models.

Recommendations

- Map TRANSFER Agent to Linked Art Actor as the primary mapping.
- Develop controlled vocabulary for TRANSFER-specific Agent roles (e.g., creator, collaborator, curator, conservator, installer, fabricator, developer) and establish guidelines on usage.
- Plan future refinement to include intermediate AgentRole class as noted in Section 4.4 Future State Considerations.
- Consider establishing guidelines for documenting both individual and organizational Agents, particularly in collaborative TBM creation and conservation contexts.

5.3.2. Partial Coverage Areas

Series

Neither Linked Art nor CIDOC-CRM provides a native mechanism for documenting artist-conceived series of related works that share thematic, formal, or procedural continuity. Linked Art introduces a Set class from outside of CIDOC-CRM to handle various grouping scenarios (Linked Art Editorial Board, n.d., *Collection*). However, the documented use cases for Set focus primarily on institutional and curatorial contexts such as museum collections, departmental collections, auction lots, dealer inventories, exhibitions, archival groupings, and sets of people sharing common features. While Set could technically accommodate series relationships through its flexible grouping mechanism, it lacks the semantic nuance needed to distinguish series from other types of groupings.

Series is a well-established concept in art historical discourse, with recognized meaning distinct from general collections or exhibitions. The Getty Art & Architecture Thesaurus provides a definition of series as "a conceptual grouping of visual arts works, literary works, or performance art created in succession by the same author, artist, or studio and intended by the creator(s) to be seen together or in succession as a cycle of works" (Getty Research Institute, n.d.). In the AAT hierarchy, series is classified under Object Groupings and Systems > object groupings > <object groupings by general context> > series (object groupings). This hierarchical placement situates series as a specific type of object grouping distinguished by its conceptual context. For TRANSFER use cases, Series differ from institutional collections or exhibition groupings in several critical ways:

- **Conceptual Intent:** Series are defined by intentional conceptual relationships, whether established by an artist during their lifetime or identified posthumously through retrospective analysis. This differs from institutional organizing principles where collections and exhibitions are typically organized by external parties for purposes such as collections management or display.
- **Thematic and Procedural Continuity:** Works within a Series may maintain specific relationships around artistic development, formal experimentation, or conceptual exploration that evolve over time. These relationships require documentation of how the series concept itself develops and how individual works contribute to or diverge from series themes.
- **Temporal Development:** Series often unfold over extended periods of an artist's practice, with works created non-consecutively as the artist returns to and refines series concepts. This differs from exhibition sets (bounded by site-specific dates) or collection sets (defined by acquisition history).
- **Identity Relationships:** Works within a series maintain their individual identities while participating in the larger series identity. This differs from auction lots (where objects are temporarily grouped for sale) or archival sets (where hierarchical containment may be more appropriate).

The generic nature of Set, designed to accommodate diverse grouping scenarios from auction lots to museum departments, means it cannot adequately capture the semantic nuance that distinguishes series from institutional collections or exhibition groupings without substantial additional modeling development. Moreover, using Set for Series would conflate fundamentally different types of groupings (institutional, curatorial, temporal, conceptual) within the same class, potentially creating ambiguity in how systems interpret and display series relationships.

PREMIS Intellectual Entity can similarly represent conceptual groupings and could accommodate Series relationships: "Intellectual Entities occur at all levels of aggregation from collections down to individual embedded images. . . PREMIS supports this modeling through flexible relationship subtypes for capturing their relationships to other objects." (PREMIS Editorial Committee, 2015, p. 11). However, like Linked Art's Set, it lacks explicit semantics for the unique characteristics of artistic series.

Recommendations

- Map TRANSFER Series to Linked Art Set with classified_as aat:300027349 (Series).
- Consider proposing a new Series class to the Linked Art community. If the Series class is approved and implemented in Linked Art, map TRANSFER Series to Linked Art Series.
- Consider establishing guidelines for determining when groups of works constitute a formal Series versus other types of relationships.

Edition

Linked Art, CIDOC-CRM, nor PREMIS provides adequate mechanisms for documenting editions as understood in collection and distribution contexts.Edition is a well-established concept across art documentation standards. Despite this gap in semantics for Linked Art and CIDOC-CRM, edition is a well-established concept across art documentation standards. VRA Core 4.0 includes stateEdition as an element in their core schema, defining it as "the identifying number and/or name assigned to the state or edition of a work that exists in more than one form and the placement of that work in the context of prior or later issuances of multiples of the same work" (Visual Resource Association, 2007). Similarly, CDWA states that an edition "identifies a specific work in the context of a group issued at the same time, or defines an issuance of a work in relation to previous and subsequent editions" (CDWA, 2024). While Linked Art offers a generic Set class for grouping objects, it explicitly states "we don't have 'Edition' as an entity to associate information with" when mapping to Categories for the Description of Works of Art (Linked Art Editorial Board, n.d., *CDWA*). This absence represents a significant gap that extends beyond time-based media to affect documentation of prints, photographs, sculptures, and other multiples across all artistic practices.

Critical edition semantics missing from current standards include:

- Edition numbering (e.g., 3/5, 1/1, artist's proof)
- Rights and ownership structures specific to individual Editions
- Pricing and valuation associated with Edition units
- Distribution terms and conditions
- Relationships between Editions of the same Work (e.g., how a limited edition relates to artist's proofs)
- Edition-specific provenance independent of individual Variant-level provenance

The absence of Edition modeling creates significant challenges for artist-driven archives, galleries, and collectors who must track distribution, ownership, and value at the edition level rather than treating each instantiation as a unique work.

Recommendations

- Map Edition to Set with classified_as aat:300121294 (Edition).
- Consider proposing a new Edition class to the Linked Art community. If the Edition class is approved and implemented in Linked Art, map TRANSFER Edition to Linked Art Edition.

Variant

Linked Art maintains separate HumanMadeObject and DigitalObject classes, reflecting CIDOC-CRM's distinction between E22 Human-Made Object (physical) and E73 Information Object (digital). This binary classification presents challenges for TBM works that inherently exist across both physical and digital modalities simultaneously, such as projection-based installations that require both digital files and physical display infrastructure.

PREMIS provides superior support for hybrid materiality through its Representation class, defined as "a single instance of an Intellectual Entity held in a preservation repository" (PREMIS Editorial Committee, 2015, p. 12). This approach accommodates the complex material reality of TBM works without forcing artificial distinctions between “human-made” and digital components. A salient example from PREMIS documentation illustrates this flexibility: "'Statue of a horse' might be a part of an article consisting of that TIFF image and a file of SGML-encoded text. If the repository created a JPEG2000 version of the TIFF, it would hold two Representations of the article: the TIFF and the SGML files would make up one representation, while the JPEG2000 and the SGML files would make up another representation" (PREMIS Editorial Committee, 2015, p. 12).

Recommendations

- Map TRANSFER Variant to Linked Art HumanMadeObject with classified_as using local term "Variant."
- Consider proposing an AAT term for Variant to replace local vocabulary.
- Consider proposing a new Variant class to the Linked Art community. If the Variant class is approved and implemented in Linked Art, map TRANSFER Variant to Linked Art Edition.

Iteration

Iteration maps partially to Linked Art's Exhibition activity, which provides a strong technical foundation through the underlying Activity class for documenting temporal events with nuance. However, a significant semantic gap exists between "exhibition" and the established concept of "iteration" in TBM conservation practice.

In TBM conservation discourse, the term "iteration" has a specific meaning rooted in the allographicity of time-based media work, where the composition and execution are separate stages. This concept extends to conceptual installation art, which "necessitates an iterative approach if it is to persist as an artwork beyond the ephemeral, and each installation is unavoidably a new translation of its concept" (Bloser, 2018). Performance art similarly operates through iteration, as "each iteration of a live performance is unique and unrepeatable," with the work existing only through its repeated instantiations across different contexts (Coogan, 2015). Because time-based works only exist when installed and each viewing represents a unique instantiation, "a new viewing of a time-based work can only be an iteration" (Phillips, 2015). Simply put, an iteration represents a time- and site-specific execution that may occur outside traditional exhibition contexts (e.g., studio testing, private installations, artist-directed presentations), making Exhibition semantically incomplete for representing Iteration.

PREMIS Event provides an alternative mapping that captures preservation actions and temporal specificity, but lacks the artistic, exhibition, spatial, and decision-making dimensions that define iterations in TBM conservation practice. PREMIS focuses primarily on technical preservation activities (e.g., migration, validation, fixity checking) rather than the experiential, collaborative, and interpretive aspects that iteration reports document. The term "event" in PREMIS encompasses any action performed on a digital object without the nuanced understanding of how an artwork is realized in a specific time and place that is fundamental to the iteration concept.

Given these limitations in both Linked Art and PREMIS for fully capturing Iteration, TRANSFER should pursue a dual-track approach. In the short term, mapping Iteration to Linked Art's Activity class provides immediate access to robust event-based modeling capabilities while acknowledging the semantic gap. Concurrently, TRANSFER should propose a new Iteration class to the Linked Art community that formally codifies the specific semantics of time- and site-specific executions in TBM, conceptual installation art, and performance art contexts. If the Iteration class is approved and implemented in Linked Art, the mapping can then be refined from the broader Activity subclass relationship to a more precise owl:equivalentClass assertion between TRANSFER Iteration and Linked Art Iteration.

Recommendations

- Map TRANSFER Iteration to Linked Art Activity as an interim implementation approach.
- Propose a new Iteration class to the Linked Art community. If the Iteration class is approved and implemented in Linked Art, map TRANSFER Iteration to Linked Art Iteration.

Version

Neither Linked Art nor CIDOC-CRM provides native mechanisms for versioning technical or procedural updates to artworks. This represents a significant gap for TBM conservation documentation, where tracking software updates, hardware substitutions, encoding changes, and technical environment modifications are essential for preservation planning. TBM works frequently undergo technical evolution that changes how they function or are experienced without fundamentally altering their conceptual identity or modality. A projection-based work might be updated from DVD playback to file-based playback, require new codec versions as operating systems evolve, or necessitate replacement display equipment as original hardware becomes obsolete. These changes are distinct from creating new Variants, which represent different modalities like projection versus web-based, or documenting new Iterations, which represent time- and site-specific installations. Instead, they represent serial technical development where each Version builds upon and potentially replaces previous technical implementations.

PREMIS offers comprehensive support through two distinct versioning mechanisms that provide valuable granularity for preservation contexts. The agentVersion property is "used to express the version of software Agents" (PREMIS Editorial Committee, 2015, p. 5). For example, agentVersion might capture "Adobe Photoshop CS2" or "QuickTime 7.6.2" when documenting the specific software used to create or render digital content. The environmentalVersion property, by contrast, identifies the version of an environmental stack (e.g., hardware architecture, hardware peripheral, operating system, software application, software driver, software library). This distinction enables precise documentation of both the direct creative tools (agents) and the broader technical infrastructure (environment) required for artwork execution. While these PREMIS properties could theoretically be modeled as subclasses of Version, this approach would fragment version documentation across multiple entity types and lacks a direct mapping to a class-level concept that represents the Version itself as a coherent unit of change.

Recommendations

- Map TRANSFER Version to Linked Art InformationObject with classified_as aat:300121393 (Version).
- Propose a new Version class to the Linked Art community. If the Version class is approved and implemented in Linked Art, map TRANSFER Version to Linked Art Version.
- Consider proposing guidelines for naming or numbering conventions for Version identifiers that facilitate serial tracking and dependency management.

Component

Linked Art models components as part-whole relationships using the part_of property (CIDOC-CRM P46 is composed of) rather than as a distinct class. This property-based approach allows components to be represented through HumanMadeObject or DigitalObject entities connected to parent objects, but it fundamentally limits the expressive power needed for TBM conservation documentation.

The position of this analysis is that Component must be modeled as a first-class entity with its own identity, properties, and relationships, rather than merely as a property connecting parts to wholes. This position aligns with established collection management practices. For example, Gallery Systems' The Museum System provides dedicated component management functionality that treats components as distinct objects with their own records, relationships, and lifecycle tracking (Gallery Systems, n.d.). Tate's research on collection management for complex artworks similarly addresses the need to register and document components as independent entities that enable tracking of conservation requirements, technical specifications, and replacement scenarios across multiple works. This institutional practice extends beyond TBM to complex sculptures, installations, and mixed-media works across all artistic practices (Huyton, 2019).

Components may have identities and lifecycles independent of any single Work. A media player, display monitor, or software package may be used across multiple Variants, replaced due to obsolescence, maintained independently, or documented before assignment to any specific Work. Components also participate in complex relationships beyond simple part-whole hierarchies. They may be required by multiple Variants simultaneously, depend on other Components for functionality, be substitutable by alternatives, undergo maintenance independent of installations, or be specified at the Version level while used at the Iteration level. These many-to-many relationships across different entity types cannot be adequately modeled through properties alone.

PREMIS provides strong support for digital components through its File, Object, and Environment entities. However, PREMIS lacks adequate semantics for physical sculptural elements, fabricated display surfaces, and other components with aesthetic significance beyond their technical function. PREMIS does not capture material properties, dimensions, spatial relationships, or the artistic/conservation importance of physical components, which are all critical for comprehensive TBM documentation.

Recommendations

- Map TRANSFER Component to Linked Art HumanMadeObject or DigitalObject with classified_as aat:300241490 (Component)
- Propose a new Component class to the Linked Art community. If the Component class is approved and implemented in Linked Art, map TRANSFER Component to Linked Art Component.
- Consider developing specialized properties for component documentation. Examples:
    - technicalSpecifications (structured or free-text technical details)
    - dependsOn (relationships to other required components)
    - compatibleWith (version or technical compatibility information)
    - substitutableBy (acceptable replacement components)
    - expectedLifespan (for maintenance planning)
    - obsolescenceStatus (current availability and replacement considerations)
- Establish clear guidelines for determining appropriate granularity of component documentation (e.g., when to document individual cables versus cable assemblies).

7. Appendices

7.1. Mermaid Diagram Syntax

The conceptual data model diagram (Section 4.3) was created using Mermaid, an open-source diagramming syntax that generates diagrams from text descriptions. The syntax below can be rendered using any Mermaid-compatible tool, including:

- **Online editors:** Mermaid Live Editor for immediate rendering and export
- **Documentation platforms:** GitHub, GitLab, and Notion support native Mermaid rendering
- **Development tools:** VS Code extensions, Markdown editors with Mermaid support
- **Static site generators:** Jekyll, Hugo, and others with Mermaid plugins

To use this syntax, copy the code below into any Mermaid renderer to generate an interactive version of the conceptual model diagram.

```
---
config:
  layout: elk
---
erDiagram
    Series ||--o{ Work : includes
    Series }o--o{ Series : references
    Work ||--o{ Edition : hasEdition
    Work }o--o{ Work : references
    Edition ||--o{ Variant : hasVariant
    Variant }|--o{ Version : hasVersion
    Variant }|--o{ Iteration : hasIteration
    Variant }|--o{ Component : requires
    Version }o--o{ Version : derivesFrom
    Version }|--o{ Iteration : usedIn
    Version }|--o{ Component : specifies
    Iteration }o--o{ Iteration : iteratesOn
    Iteration }|--o{ Component : uses
    Component }o--o{ Component : dependsOn
    Activity ||--o{ Series : conceives
    Activity ||--o{ Work : creates
    Activity ||--o{ Edition : produces
    Activity }|--o{ Variant : defines
    Activity }|--o{ Version : versions
    Activity ||--|| Iteration : isA
    Activity }|--o{ Component : maintains
    Agent }o--o{ Activity : carriesOut
```

7.2. Relevant Linked Art Model Documentation

- Abstract Works → https://linked.art/api/1.0/endpoint/abstract_work/

- CIDOC-CRM Class Analysis → https://linked.art/model/profile/class_analysis/
- Digital Content → https://linked.art/model/digital/
- Digital Objects → https://linked.art/api/1.0/endpoint/digital_object/
- Exhibitions → https://linked.art/model/exhibition/
- Objects → https://linked.art/model/object/
- Rights → https://linked.art/model/object/rights/
- Sets → https://linked.art/api/1.0/endpoint/set/

7.3. Relevant CIDOC CRM Class Definitions

E7 Activity

Subclass of:	E5 Event
Superclass of:	E8 Acquisition E9 Move E10 Transfer of Custody E11 Modification E13 Attribute Assignment E65 Creation E66 Formation E85 Joining E86 Leaving E87 Curation Activity
Scope note:	This class comprises actions intentionally carried out by instances of E39 Actor that result in changes of state in the cultural, social, or physical systems documented. This notion includes complex, composite, and long-lasting actions such as the building of a settlement or a war, as well as simple, short-lived actions such as the opening of a door.
Examples:	• the Battle of Stalingrad (Hoyt, 1993) • the Yalta Conference (Harbutt, 2010) • my birthday celebration 28-6-1995 • the writing of "Faust" by Goethe (E65) (Williams, 2020) • the formation of the Bauhaus 1919 (E66) (Droste, 2006) • the people of Iraq giving the name ‘Quyunjig’ to the place identified by the TGN as‘7017998’ • Kira Weber working in glass art from 1984 to 1993 • Kira Weber working in oil and pastel painting from 1993
Properties:	P14 carried out by (performed): E39 Actor (P14.1 in the role of: E55 Type) P15 was influenced by (influenced): E1 CRM Entity P16 used specific object (was used for): E70 Thing (P16.1 mode of use: E55 Type) P17 was motivated by (motivated): E1 CRM Entity P19 was intended use of (was made for): E71 Human-Made Thing (P19.1 mode of use: E55 Type) P20 had specific purpose (was purpose of): E5 Event P21 had general purpose (was purpose of): E55 Type

	P32 used general technique (was technique of): E55 Type P33 used specific technique (was used by): E29 Design or Procedure P125 used object of type (was type of object used in): E55 Type P134 continued (was continued by): E7 Activity

E28 Conceptual Object

Subclass of:	E71 Human-Made Thing
Superclass of:	E55 Type E89 Propositional Object E90 Symbolic Object
Scope note:	This class comprises non-material products of our minds and other human produced data that have become objects of a discourse about their identity, circumstances of creation, or historical implication. The production of such information might have been supported by the use of technical devices such as cameras or computers. Characteristically, instances of this class are created, invented or thought by someone, and then may be documented or communicated between persons. Instances of E28 Conceptual Object have the ability to exist on more than one particular carrier at the same time, such as paper, electronic signals, marks, audio media, paintings, photos, human memories, etc. They cannot be destroyed. They exist as long as they can be found on at least one carrier or in at least one human memory. Their existence ends when the last carrier and the last memory are lost.
Examples:	● Beethoven’s "Ode an die Freude" (Ode to Joy) (E73) (Kershaw, 1999) ● the definition of "ontology" in the Oxford English Dictionary (E73) (Oxford University Press, 1989) ● the knowledge about the victory at Marathon carried by the famous runner (E89) (Lagos & Karyanos, 2020)
Properties:	N/A

E39 Actor

Subclass of:	E77 Persistent Item
Superclass of:	E21 Person E74 Group
Scope note:	This class comprises people, either individually or in groups, who have the potential to perform intentional actions of kinds for which they can be held responsible.
Examples:	● London and Continental Railways (E74) ● the Governor of the Bank of England in 1975 (E21)

	• Sir Ian McKellen (E21) (Gibson, 1986)
Properties:	P74 has current or former residence (is current or former residence of): E53 Place P75 possesses (is possessed by): E30 Right P76 has contact point (provides access to): E41 Appellation

E71 Human-Made Thing

Subclass of:	E70 Thing
Superclass of:	E24 Physical Human-Made Thing E28 Conceptual Object
Scope note:	This class comprises discrete, identifiable human-made items that are documented as single units. These items are either intellectual products or human-made physical things, and are characterized by relative stability. They may, for instance, have a solid physical form, an electronic encoding, or they may be logical concepts or structures.
Examples:	• Beethoven's 5th Symphony (E73) (Lockwood, 2015) • Michelangelo's David (E22) (Paoletti and Bagemihl, 2015) • Einstein's Theory of General Relativity (E89) (Hartle, 2003) • the taxon 'Fringilla coelebs Linnaeus,1758' (E55) (Sinkevicius and Narusevicius, 2002)
Properties:	P102 has title (is title of): E35 Title (P102.1 has type: E55 Type) P103 was intended for (was intention of): E55 Type

E89 Propositional Object

Subclass of:	E28 Conceptual Object
Superclass of:	E73 Information Object E30 Right
Scope note:	This class comprises immaterial items, including but not limited to stories, plots, procedural prescriptions, algorithms, laws of physics or images that are, or represent in some sense, sets of propositions about real or imaginary things and that are documented as single units or serve as topic of discourse. This class also comprises items that are "about" something in the sense of a subject. In the wider sense, this class includes expressions of psychological value such as non-figural art and musical themes. However, conceptual items such as types and classes are not instances of E89 Propositional Object. This should not be confused with the definition of a type, which is indeed an instance of E89 Propositional Object.
Examples:	• Maxwell's Equations (Ball, 1962) • the ideational contents of Aristotle's book entitled 'Metaphysics' as rendered in

	the Greek ● texts translated in Oxford edition ● the underlying prototype of any "no-smoking" sign (E36) ● the common ideas of the plots of the movie "The Seven Samurai" by Akira Kurosawa and the ● movie "The Magnificent Seven" by John Sturges (Mellen, 2002) ● the image content of the photo of the Allied Leaders at Yalta published by UPI, 1945 (E36) ● the character "Little Red Riding Hood", variants of which appear amongst others in Grimm brothers’ "Rotkäppchen", other oral fairy tales and the film "Hoodwinked" ● the place "Havnor" as invented by Ursula K. Le Guin for her "Earthsea" book series, the related maps and appearing in derivative works based on these novels
Properties:	P67 refers to (is referred to by): E1 CRM Entity (P67.1 has type: E55 Type) P129 is about (is subject of): E1 CRM Entity P148 has component (is component of): E89 Propositional Object

7.4. References

Alexopoulos, P. (2020). *Semantic modeling for data: Avoiding pitfalls and breaking dilemmas.* O'Reilly Media.

Bloser, J. (2018, July 27). Authenticity in installation art. *VoCA: Voices in Contemporary Art*. https://voca.network/blog/2018/07/27/authenticity-in-installation-art/.

CIDOC CRM Special Interest Group (2024). *Definition of the CIDOC Conceptual Reference Model version 7.1.3*. https://cidoc-crm.org/sites/default/files/cidoc_crm_version_7.1.3.pdf.

Coogan, A. (2015, November 11). *What is performance art?* https://imma.ie/magazine/what-is-performance-art/.

Gallery Systems (n.d.). Bridging the object-component divide. Retrieved from https://www.gallerysystems.com/bridging-the-object-component-divide/.

Getty Research Institute (n.d.). Series (object groupings). In *Art & architecture thesaurus*. http://vocab.getty.edu/page/aat/300027349.

Getty Research Institute. (2024). Edition. In *Categories for the Description of Works of Art* (E. Benoff, M. Baca, & P. Harpring, Eds.). J. Paul Getty Trust. https://www.getty.edu/publications/categories-description-works-art/categories/object-architecture-group/10/.

Hitzler, P., Gangemi, A., Janowicz, K., Krisnadhi, A. A., and Presutti, V. (2017). Towards a simple but useful ontology design pattern representation language. In Blomqvist, E., Corcho, Ó., Horridge, M., Carral, D., and Hoekstra, R., editors, *Proceedings of the 8th Workshop on Ontology Design and Patterns (WOP 2017) co-located with the 16th International Semantic Web Conference (ISWC 2017), Vienna, Austria*, October 21, 2017., volume 2043 of *CEUR Workshop Proceedings*. https://people.geog.ucsb.edu/~jano/wop17.pdf.

Hyvönen, E. (2012). *Publishing and using cultural heritage linked data on the Semantic Web*. San Rafael, CA: Morgan & Claypool.

Huyton, S. (2019). Registration and collection management. In *Reshaping the Collectible: When Artworks Live in the Museum*. Tate. https://www.tate.org.uk/research/reshaping-the-collectible/research-approach-registration-collection-management.

Linked Art Editorial Board (n.d.). *CDWA*. https://linked.art/cookbook/mappings/cdwa/#10-edition.

Linked Art Editorial Board (n.d.). *Collection*. https://linked.art/model/collection/.

Linked Art Editorial Board (n.d.). *Profile*. https://linked.art/model/profile/.

Newbury, D. (2018). LOUD: Linked Open Usable Data and linked.art. In *CIDOC 2018: Enrichment and Augmentation of Cultural Heritage – Proceedings of the 2018 Annual Conference of CIDOC*, Heraklion, Crete, Greece, September 29 – October 5, 2018. ICOM. https://cidoc.mini.icom.museum/wp-content/uploads/sites/6/2021/03/CIDOC2018_paper_153.pdf.

Noy, N. F., & McGuinness, D. L. (2001). *Ontology development 101: A guide to creating your first ontology*. https://protege.stanford.edu/publications/ontology_development/ontology101.pdf.

Phillips, J. (2015). Reporting iterations: A documentation model for time-based media art. In L. A. Matos, R. Macedo, & G. Heydenreich (Eds.), *Revista de História da Arte: Performing Documentation in the Conservation of Contemporary Art*, 4, 169-177.

PREMIS Editorial Committee (2015). *PREMIS Data Dictionary for preservation metadata version 3.0*. https://www.loc.gov/standards/premis/v3/premis-3-0-final.pdf.

Taylor, A. (2015). *Semantics for dummies, MarkLogic special edition*. Hoboken, NJ: John Wiley & Sons, Inc.

Triplestore. (n.d.). In *Wikipedia*. Retrieved from https://en.wikipedia.org/wiki/Triplestore.

Uschold, M., & Gruninger, M. (1996). Ontologies: Principles, methods and applications. *Knowledge Engineering Review* 11(2). https://www.aiai.ed.ac.uk/publications/documents/1996/96-ker-intro-ontologies.pdf.

Visual Resources Association. (2007). *VRA Core 4.0 element description*. Library of Congress. https://www.loc.gov/standards/vracore/VRA_Core4_Element_Description.pdf.

World Wide Web Consortium (W3C). (n.d.). *Vocabularies*. Retrieved from https://www.w3.org/standards/semanticweb/ontology.