Technical Report

ISO/TR 24332

First edition
2025-01

Information and documentation — Blockchain and distributed ledger technology (DLT) in relation to authoritative records, records systems and records management

Information et documentation — Chaînes de blocs et dispositifs d'enregistrement électronique (DEE) partagé en lien avec les enregistrements officiels, les systèmes d'enregistrement et la gestion des enregistrements

Reference number
ISO/TR 24332:2025(en)

ISO copyright office
CP 401 • Ch. de Blandonnet 8
CH-1214 Vernier, Geneva
Phone: +41 22 749 01 11
Email: copyright@iso.org
Website: www.iso.org

Published in Switzerland

Contents

Page

Foreword

ISO (the International Organization for Standardization) is a worldwide federation of national standards bodies (ISO member bodies). The work of preparing International Standards is normally carried out through ISO technical committees. Each member body interested in a subject for which a technical committee has been established has the right to be represented on that committee. International organizations, governmental and non-governmental, in liaison with ISO, also take part in the work. ISO collaborates closely with the International Electrotechnical Commission (IEC) on all matters of electrotechnical standardization.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular, the different approval criteria needed for the different types of ISO document should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives).

ISO draws attention to the possibility that the implementation of this document may involve the use of (a) patent(s). ISO takes no position concerning the evidence, validity or applicability of any claimed patent rights in respect thereof. As of the date of publication of this document, ISO had not received notice of (a) patent(s) which may be required to implement this document. However, implementers are cautioned that this may not represent the latest information, which may be obtained from the patent database available at www.iso.org/patents. ISO shall not be held responsible for identifying any or all such patent rights.

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation of the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT), see www.iso.org/iso/foreword.html.

This document was prepared by Technical Committee ISO/TC 46, *Information and documentation*, Subcommittee SC 11, *Archives/records management*, in collaboration Technical Committee ISO/TC 308, *Blockchain and distributed ledger technologies*.

Any feedback or questions on this document should be directed to the user's national standards body. A complete listing of these bodies can be found at www.iso.org/members.html.

Introduction

Distributed ledger technology (DLT), including blockchain technology, is expected to be widely adopted for business and governance purposes. The viability of this technology is already established in many contexts, and DLT solutions can potentially be used in any industry, sector or context.

Information systems used for business and governance can create, receive and store records. DLT solutions are no different. There can be records in these solutions that need to be managed in compliance with existing legal, regulatory, business, societal and other requirements. Also, DLT solutions or their constituent parts have potential to be designed to manage records.

The need for the analysis of DLT from a records management point of view results from the specific characteristics of this technology (e.g. distributed and decentralized nature, immutability, use of consensus and use of smart contracts) and some of its modes of application (e.g. including the possibility of there being no designated owner, distributed governance, transborder use, and different trust assumptions). The specific characteristics of DLT can both facilitate records management (e.g. maintenance of integrity) and result in difficult records management and legal challenges [e.g. possible absence of a designated authoritative copy of a record, difficulties in disclosing records to authorities and courts including e-discovery, difficulties managing retention and disposition, and challenges managing personally identifiable information (PII) protection].

This document provides a foundation for common understanding among records managers, DLT system designers/developers and related professionals and can inform the development of future standards concerning DLT and records management. This document does not presume in depth knowledge of records management principles or DLT.

Information and documentation — Blockchain and distributed ledger technology (DLT) in relation to authoritative records, records systems and records management

1 Scope

This document analyses challenges, considerations, and potential benefits of blockchain and distributed ledger technology (DLT) in relation to records management standards and related standards for systems that:

— create records that are required to be authoritative records;

— can be used as records systems; or

— can be used for records management, including records controls.

The target audience of this document includes records managers and allied professionals, IT professionals and application developers, legal and compliance professionals, researchers, educators and other interested parties.

2 Normative references

The following documents are referred to in the text in such a way that some or all of their content constitutes requirements of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ISO 22739, *Blockchain and distributed ledger technologies — Vocabulary*

ISO 30300, *Information and documentation — Records management — Core concepts and vocabulary*

3 Terms and definitions

For the purposes of this document, the terms and definitions given in ISO 22739 and ISO 30300 apply.

ISO and IEC maintain terminological databases for use in standardization at the following addresses:

— ISO Online browsing platform: available at https://www.iso.org/obp

— IEC Electropedia: available at https://www.electropedia.org/

4 Overview of records management and distributed ledger technology (DLT)

4.1 General

Any governance or business activity heavily relies on authoritative records, rather than on just any information or data. Authoritative records are essential for decision-making, protection of rights, transparency, accountability and memory. As soon as DLT solutions are used for business and governance, they can create, receive and keep records that are subject to legal, regulatory, business and other requirements, and can also have long-term or historical value.

To manage these records, one can rely on extensive body of knowledge and practical experience reflected in ISO records management standards.

Records management is the discipline responsible for the efficient and systematic governance of records using records processes, records controls and records systems. Understanding records as information created or received and maintained as evidence of conducting business, records processes are a set of activities for creating, capturing and managing authoritative records. These activities are supported by records controls, such as business classification schemes or metadata schemas, and are performed in records systems or across an organization.

DLT, which includes blockchain technology, enables the operation and use of distributed ledgers containing transaction records that are intended to be final, definitive and immutable.

To help understand this document, this clause introduces overviews of the records management and DLT based on the following International Standards.

— ISO 30300: provides the most relevant definitions and concepts diagrams related to the concepts used in the records management domain;

— ISO 30301: specifies requirements to be met by a management system for records;

— ISO 15489-1: establishes the core concepts and principles for the creation, capture and management of records;

— ISO 23257: specifies a reference architecture for DLT;

— ISO 22739: specifies vocabulary for DLT.

These documents can be consulted for more detailed advice on aspects of managing records or DLT.

4.2 Overview of records management concepts and principles

4.2.1 Concepts

Records are both evidence of business processes, activities and transactions and information assets. Any set of information, regardless of its structure or form, can be managed as a record. The creation, capture and management of records are integral parts of conducting business, in any context. Records document individual events or transactions or can form aggregations that have been designed to document business processes, activities or functions.

Evidence is understood as information that can be used either by itself or in conjunction with other information, to establish proof about an event or action. Evidence is not limited to the legal sense of the term. Records that possess the characteristics of authenticity, reliability, integrity and useability are considered authoritative evidence. Records that have these characteristics are called authoritative records.

Metadata for records is data describing the context, content and structure of records, as well as their management over time (see ISO 23081).

Records that do not possess such metadata are generally not considered authoritative.

Decisions regarding the creation, capture and management of records are based on the analysis and risk assessment of business functions, processes and activities, in their business, legal, regulatory and societal contexts. The analysis process is called appraisal (see ISO/TR 21946).

4.2.2 Principles for managing records

Managing records encompasses the following:

— establishing management systems for records

— creating and capturing records to meet requirements for evidence of business activity;

— taking appropriate action to maintain and protect their authenticity, reliability, integrity and useability as their business context and requirements for their management change over time.

A management system for records is a set of interrelated elements used to direct and control an organization with regard to records. Elements include leadership, policy, planning, resources and other supports, operations, performance evaluation and continual improvement.

Records management operations are supported by processes for creating records, capturing records, classification and indexing, access control, storing records, use and reuse, migration or conversion and disposition (retention, destruction or transfer) of records. These records processes rely on records controls which are instruments designed specifically to help in their performance such as metadata schemas for records, business classification schemes, access and permissions rules and disposition authorities.

The management of records is supported by records systems which are information systems that are designed specifically to manage records, or that are designed for other business processes that are adapted to support the management of records.

Continuous monitoring and evaluation are essential to ensure that records management practices remain effective and aligned with evolving business needs.

4.3 Overview of distributed ledger technology (DLT) and blockchain

4.3.1 General

Ledgers underlie accounting, commerce, taxation, and the orderly conduct of economies. Historically, ledger technologies have included physical tokens, tally sticks, double-entry books, and centralized computerized information systems. Blockchain and DLT are a new kind of computerized ledger technology, where ledgers are not just distributed (in their physical structure) but can also be decentralized (in their control structure). A blockchain system is one type of DLT system, but some DLT systems are not blockchain systems. In the remainder of this document, DLT includes blockchain technology, and only distinguishes them when required.

Blockchain technology was introduced by the Bitcoin platform, which demonstrated a solution to the long-standing challenge of how to enable digital cash. Digital cash, like traditional physical cash (and unlike bank deposits) can be directly controlled by its owner, but like bank deposits (and unlike physical cash) can be transferred to remote parties globally. A challenge for digital cash systems is to ensure that every unit of digital cash has no more than one owner at a time even without a centralized authority (often referred to as the "double spending problem"). Solving this is difficult because information goods are not inherently exclusionary. The Bitcoin platform realized digital cash as the Bitcoin cryptocurrency, and Bitcoin's ledger was defined by a blockchain which represented all transactions of transfers of Bitcoin (and associated data) in a single globally visible list of transactions.

DLT can account not just for money, but also for other kinds of assets. DLT systems after Bitcoin have expanded the capability of their ledgers to be able not just to represent cryptocurrencies, but also other kinds of digital assets, data, and programs called "smart contracts". Smart contracts are recorded in a DLT system, and their results of execution are also recorded on the ledger. DLT systems can, either through validation in the platform or through smart contract execution, enforce integrity conditions for digital assets, data, and smart contracts on their ledgers. Consequently, just like modern centralized databases, modern DLT systems can be used as general-purpose data storage, computation, and communication components in information systems. DLT systems typically have some limitations compared to centralized database systems (such as for performance efficiency and confidentiality) but can have some advantages (such as for availability and integrity).

4.3.2 Overview of the distributed ledger technology (DLT) reference architecture

A reference architecture is a common generic model for a class of systems. The reference architecture for DLT systems describes both the internal architecture of underlying DLT platforms, and the related non-DLT systems that all together implement solutions for specific use cases. The reference architecture standard ISO 23257 describes a range of overall DLT concepts and identifies important cross-cutting aspects for DLT

systems. In the design of software systems, architectural decisions are important in addressing these cross-cutting aspects, which include qualities such as security, and performance efficiency, and other aspects such as identity, governance, and management of DLT systems. The reference architecture standard ISO 23257 outlines how the decentralized nature and typical structure of DLT systems impacts the achievement of requirements for these cross-cutting aspects.

An important part of a reference architecture is the set of architectural views. Each view models a system relative to a specific set of concerns. The DLT reference architecture provides three views.

— User view – the roles and responsibilities associated with DLT systems, including users, providers, developers, administrators, governors, and auditors.

— Functional view – the functionalities within and provided by DLT systems. These are grouped by coarse "layers": Infrastructure Layer, DLT Platform Layer, API Layer, Non-DLT Systems, User Layer, and Cross-Layer Functions. Non-DLT systems include DLT oracles that provide a gateway for external data to a DLT system, non-DLT applications that can interoperate with the DLT system, and off-ledger data that can relate to the DLT system.

— System view – the structural elements within and connected to DLT systems. A DLT system is implemented by a network of DLT nodes, each of which runs the DLT platform. The platform provides API interfaces to users, and other interfaces to external non-DLT systems and other DLT systems. Within a DLT platform, there are elements including the ledger, transaction and consensus mechanisms, smart contracts, and cryptographic services. Spanning the whole DLT system are other elements such as infrastructure services, and other cross-layer elements for development, management and operations, security, and governance and compliance.

4.3.3 Different types of distributed ledger technology (DLT) systems

There are different kinds of DLT systems. They differ in five important aspects: access for use, authorization, ledger structure, smart contract capability, and consensus mechanism.

DLT systems can be public, in which case access for use is available to all, or private, in which case access for use is restricted to a limited group of participants. Privacy is not guaranteed even in the case of private DLT systems, because all the DLT nodes participating in the consensus mechanism for a transaction will typically have access to the information in that transaction. Private DLT systems tend to be smaller, and can have well-known and more trustworthy DLT nodes, and therefore often have better performance than public DLT systems. However, public DLT systems can provide high levels of transparency and integrity through wide public participation and oversight.

A DLT system can be permissionless, in which case authorization is not required to perform activities in the system, or can be permissioned, in which case authorization is required to perform at least some activities.

The ledger structure of a blockchain system is a linked chain: a single global list of transactions, grouped into cryptographically linked blocks, each of which contains a list of transactions. However, other kinds of DLT systems can have different ledger structures, which can help to improve concurrency and performance. Some DLT systems fragment the ledger into multiple shards, to improve scalability. In other DLT systems, instead of there being a global ledger, there are many small ledgers, shared just between parties of interest to their transactions.

DLT systems can vary in their smart contract capability. For example, some smart contract languages are "Turing-complete" and so are in principle as expressive as every other programming language. In practice, smart contract execution is usually highly resource-constrained, so that it will complete within the time and space constraints of the DLT system's consensus mechanism. Several DLT systems use expressive but sub-Turing complete smart contract languages, so that the smart contracts are more amenable to automatic static analysis or formal verification, to provide assurance about their correctness. Some DLT systems (such as the original Bitcoin blockchain) have very limited linear scripting capabilities, and some DLT systems have no smart contract capabilities.

The consensus mechanism of a DLT system enables agreement between numerous DLT nodes about the contents of the ledger. Prior to Bitcoin, a variety of consensus mechanisms were known that allowed a small number (i.e. tens) of well-known DLT nodes to reach consensus. These mechanisms included algorithms

such as Practical Byzantine Fault Tolerance, and Raft. However, in a DLT network with an unknown but large number (i.e. thousands) of DLT nodes, those approaches do not work. Bitcoin used a mechanism called Nakamoto consensus, in which DLT nodes accept as authoritative the longest ledger seen at any time. Ledgers cannot grow arbitrarily quickly in Bitcoin because of the use of a proof-of-work mechanism: blocks in the ledger must demonstrate a solution to a cryptographic puzzle which is computationally easy to check, but computationally difficult to create. Many public DLT systems continue to use Nakamoto consensus, in combination with proof-of-work, or with other approaches such as proof-of-stake. A limitation of Nakamoto consensus is that it does not provide conventional transaction properties. In conventional transaction processing, when a transaction is committed it is final, and cannot be reversed (although a reversing transaction can be subsequently committed). In Nakamoto consensus, there is only long-run probabilistic finality. At any one time, each DLT node will have their own independent view of the longest (and so, authoritative) ledger. If a DLT node is presented with a new longer but different ledger it will change what it reports as the authoritative new ledger. In practice, DLT users can reduce the likelihood of being exposed to this issue to any low-enough risk by waiting for a sufficiently long time.

4.3.4 Distributed ledger technology (DLT) use cases

DLT is a general-purpose ledger technology that can in principle be used in any sector or industry domain (see ISO/TR 3242). To satisfy the requirements of any specific use case, the design of a DLT solution will need to accommodate the limitations of the DLT system and leverage its strengths.

As demonstrated by the Bitcoin blockchain, DLT systems can support cryptocurrencies, or in general digital assets. Tokens constitute an important category of digital assets. They can represent other digital assets such as digital art or access rights, or they can represent physical assets such as museum objects. Tokens can have intrinsic value within an ecosystem or can have extrinsic value by being exchangeable for other valuable digital or physical assets. Tokens can also be used for other purposes, such as for tracking resource utilization.

Because they aim to provide a verifiable ledger, DLT systems can support efficient and trustworthy reconciliation processes and can provide coordination about mutual status and data between different individuals, businesses, or governments. These capabilities can be valuable in a wide range of industry sectors, including finance, insurance, healthcare, and supply chain management. DLT systems can enhance transparency and traceability, thereby reducing fraud and improving efficiency.

5 Distributed ledger technology (DLT) and authoritative records

5.1 General

A ledger is a long-established concept used in business and technology. Traditionally, a ledger is an information store, such as a book, containing accounts to which debits and credits are posted from books of original entry.

In the context of DLTs, a distributed ledger is a ledger that is shared across a set of DLT nodes and synchronized between the DLT nodes using a consensus mechanism. Ledger records are records containing transaction records, hash values of transaction records or references to transaction records (e.g. cryptographic links) recorded on a distributed ledger.

It is easier to use distributed ledgers for records management purposes if they are authoritative and possess the characteristics of records as described in 4.2.1. A person or organization can place confidence in an authoritative distributed ledger and rely upon it to act (i.e. to trust it).

Appraisal from the point of view of records management informs the design process for business systems that deal with records and involves the evaluation of business activities to determine which records need to be created and captured, and how long the records need to be kept. In the design of DLT solutions, designers can undertake appraisal in order to determine whether records should be created and kept on or off-ledger. Decisions about whether records and associated metadata are created and stored on or off-ledger can affect their authenticity, reliability, integrity, and useability.

5.2 On-ledger records

On-ledger records are records that are created or received, located, performed, or run inside a distributed ledger. Ledger records can contain transaction records, hash values of transaction records, or references to transaction records recorded on a distributed ledger. Ledger records can include smart contracts, which are computer programs stored in a DLT system, and include the recorded outcome of the execution of the program. Note that a smart contract can represent terms in a contract in law and create a legally enforceable obligation under the legislation of an applicable jurisdiction.

DLT systems are intended to ensure the integrity of ledger records, i.e. that the ledger and its associated records are tamper-resistant and immutable. Integrity of ledger records is supported by enforcing the validation rules of the DLT platform.

Some reasons that records are created and kept on-ledger include:

— they are necessary for the proper functioning of the DLT system (e.g. the hashes of transaction records that are used to generate the Merkle root hash that forms part of the hash used to chain blocks together in a blockchain);

— to embed reference metadata about a transaction into a ledger record to capture the context and pragmatic meaning of the transaction or for purposes of records management;

— to embed a link to transaction records or related contextual information (e.g. metadata) stored off-ledger for purposes of capturing the context and pragmatic meaning of a ledger record or to link it to supporting records related to the same transaction that are stored off-ledger. Such linkages among records or contextual information can be known by different names and in archival science are referred to as the “archival bond”.

5.3 Off-ledger records

Off-ledger records are records that are related to on-ledger records but are located in data storage outside of the DLT system. They can include any number of different types of transaction records or metadata about on-ledger records. Off-ledger data is often not immutable. A DLT system can be used as a mechanism for securing off-ledger data.

Some reasons that records are created and kept off-ledger include:

— DLT system performance, including processing efficiency;

— DLT system storage constraints;

— integration of DLT system with pre-existing business systems;

— protection of privacy and confidentiality regarding parties to a transaction or the nature of a transaction;

— legal requirements.

5.4 Metadata for records

Records can be distinguished from other information assets by their role as evidence of business and by their reliance on metadata. Metadata for records are used to indicate and preserve context and apply appropriate rules for managing records.

In DLT systems, metadata for records can be embedded into transaction records or stored as part of a record on-ledger. Alternatively, on-ledger records can link to metadata for records stored off-ledger.

5.5 Relationship between distributed ledger technology (DLT) and characteristics of authoritative records

5.5.1 General

Authoritative records possess the characteristics of authenticity, reliability, integrity and useability. This subclause discusses the relationship between DLT and characteristics of authoritative records, and the effect that decisions for the design of systems using DLT can have on the authoritativeness of records, including both on-ledger and off-ledger records.

Note that judgement about the authoritativeness of records, and thus an individual's willingness to rely upon a record to act, is complex. Records systems, including those using DLT, are designed to create and manage authoritative records. Design choices for DLT systems can impact characteristics of authoritative records in records systems using DLT. The requirement that DLT systems protect records integrity, that is, create final, definitive and immutable records, is not sufficient.

For example, the event history of records can be preserved to ensure their reliability. Typically, in records systems this is done by capturing audit or event logs; however, in DLT systems, the ledger captures the transaction history.

5.5.2 Authenticity

An authentic record is a record that can be proven to be what it purports to be, to have been created or sent by the agent purported to have created or sent it, and to have been created or sent when purported.

Business rules, processes, policies and procedures, including technical procedures, are implemented to ensure authenticity.

From the records management point of view, processes to validate and verify the identity of a records creator or to authorize the actions they can take once successfully authenticated are core components in ensuring authenticity. These components support determination that a record has been created or sent by the agent purported to have created or sent it. However, some types of DLT systems, particularly those that are public and permissionless, such as Bitcoin, do not require verification of the identity of an agent as transactions are often pseudonymous or anonymous, and are authorized cryptographically.

Establishing the authenticity of a ledger record remains possible in such DLT systems by relying upon a specific identifier (e.g. a Bitcoin address involved in the transaction that led to the creation of the ledger record) not linked to a particular agent or its legal identity if it is a record of a transaction authorized by that particular specific identifier (e.g. a transfer of Bitcoin from that address to another address).

Specific identifiers are not always sufficient to determine authenticity if the record purports to be a record authorized by a particular legally identifiable agent (e.g. a transfer of Bitcoin from one legal beneficial owner to another). In such cases, it is necessary to link the specific identifiers to the corresponding legal identifiers. Techniques that can be used to help establish authenticity without requiring identity verification and authorization include:

— using specific measures (e.g. data analytics) to determine the legal identity of a records creator with certainty or a high degree of probability;

— demonstrating that the DLT system operates in such a way as to verify the identifier of a transacting party can be associated with a particular legal identity.

Records systems sometimes include special functionality to establish authenticity, such as a link between a ledger record and metadata about the context of its creation. Depending upon the application context, the absence of such functionality can affect the ledger record's value as evidence and as an asset.

5.5.3 Reliability

A reliable record is one:

— whose contents can be trusted as a complete and accurate representation of the transactions, activities or facts to which they attest;

— which can be depended upon in the course of subsequent transactions or activities.

Records are generally considered to be more reliable if they are created at the time of the event to which they relate, or soon afterwards, by individuals who have direct knowledge of the facts, or by systems routinely used to conduct the transaction.

Reliability concerns the processes to create records as well as the completeness and accuracy of those records. The reliability of records created by a DLT system is determined by how those systems have been introduced into the workflow and the procedural controls that have been implemented to ensure reliability.

In some DLT systems, the point at which a transaction record is completed is clearly defined by the rules of the system. In other systems, completeness of a transaction record is subject to variety of considerations, including:

— when digitally signed and entered into a distributed ledger;

— when validated and confirmed;

— when confirmed and updated by a sufficient number of DLT nodes to ensure validation.

Some jurisdictions have sought to clarify the process for determining reliability by introducing legislation and regulations that establish the prima facie reliability of such records and legally recognizes smart contracts and other DLT-based records.

The security of DLT systems can impact the reliability of records they contain.

Reliability of records is impacted by the context of records creation, including the operation of the DLT system creating and storing the records.

5.5.4 Integrity

Integrity of records is the quality of being complete and unaltered.

Policies and procedures for managing records specify what additions or annotations can be made to a record after it is created, under what circumstances such additions or annotations can be authorized, and who is authorized to make them.

A key motivation of DLT systems is the protection of the integrity of ledger records. Integrity is supported when data and records are tamper-resistant and tamper-evident, i.e. that the content or sequencing of the ledger records cannot be easily altered and any changes are clearly visible. Note that tamper-resistant is a more appropriate term than tamper-proof, since it can be hard to prevent all forms of tampering. Ensuring tamper-resistance of DLT records over time is supported by comprehensive security measures and upgrades.

Similarly, DLT systems are intended to create ledger records that are maintained as immutable, specifically that they have fixity or that they are final, definitive and cannot be deleted from the ledger. This property provides for greater confidence that the records can be relied upon as evidence.

The term immutable can be misleading, however, since it implies that no changes can ever be made whereas DLT systems are designed to protect against and detect any changes to ledger records. However, from the records management point of view, the concept of integrity allows for authorized changes to records during records processes such as migration, disposition or digital preservation. This results in tension with the concept of integrity as used in IT, which refers to keeping the bit-structure unaltered. Technical and legal solutions to address this tension while still respecting the properties of tamper-resistance and immutability are being developed.

5.5.5 Useability

A useable record can be located, retrieved, presented, and interpreted over time.

A useable record is connected to the business process or transaction that produced it. Maintaining linkages between records that document related transactions ensure that records are useable.

Metadata for records such as identifiers, format or storage information, support useability and is often used to instantiate and maintain linkages between records that document transactions. Metadata for records support useability by providing information that can be needed to retrieve, present and understand them.

As with other IT infrastructure, DLT systems can lack functionality to link records to the business process or transaction that produced it, and they can also lack functionality to create linkages between records that document related transactions. That is, they are not able to instantiate the linkage between records and context or interrelationships among records (the “archival bond”).

It can be challenging to guarantee that ledger records can be located, retrieved, presented and interpreted over time. Digital preservation strategies for these records have not yet been established.

6 Distributed ledger technology (DLT) and records processes

6.1 Creating records

In each use case, creation of records on DLT can be affected by specific business, legal and other requirements.

Identifying features (e.g. system-generated metadata) and integrity controls used in the process of records creation affect the ability to establish the authenticity of records. The creator of the records, how the records were created, and how accurate was the representation of what the creator intended can affect the authenticity of records. A creator is not necessarily identified.

In case of off-ledger records in which either the hash values or references, or both, are stored on a distributed ledger and whose creation process and control are beyond the scope of operation of the DLT system, the originally hashed off-ledger records need to be archived separately in a form that is unchanged and inviolate. This enables later comparison with the hash value stored on the DLT system for purposes of checking the integrity of records stored off-ledger. If the off-ledger records are altered for any purpose (e.g. digital preservation), then the existing hash value will no longer work to validate their integrity.

To ensure the authenticity of the off-ledger records, it can be an approach to link the identity of the creator of the off-ledger record and the hash value on the distributed ledger, and to trace the provenance of both the off-ledger record and the hash value on the distributed ledger over time.

6.2 Capturing records

Capturing records is a deliberate action that results in the registration of a record into a records system. Capturing records involves capturing or generating metadata for records at the point of capture, and creation of relationships between the record and other records, agents or business functions and processes.

Though DLT systems are systems that create and store records, they are not necessarily designed to capture records according to records management standards. This can cause difficulties in assuring the authenticity of records. For instance, records not linked to their context, provenance, functions or other related records. Or metadata not conforming to the metadata schemas provided by records management standards.

6.3 Records classification

Classification links records to their business context by associating them with categories in a business classification scheme.

Records classification can include the following:

— linking the record to the business being documented, at an appropriate level (for example to a function, activity or work process);

— providing linkages between individual records and aggregations, to provide a continuous record of business activity.

As business systems, DLT systems are unlikely to have functions for creating relationships between the record and other records, agents, or business. As such, it can be difficult to link records to their procedural context or to other entities logically related to them.

6.4 Access control

Records systems are designed to support the provision and restriction of access to records. Access can be managed for individual agents, groups, or roles. Access control is used for creating, reading, updating, or deleting information. From the records management point of view, access control is helpful for ensuring and maintaining integrity, authenticity, and confidentiality of records.

Access control in both business systems and dedicated records systems usually rely on identification of users. Identifiers typically identify a natural person, a legal entity, a thing or a process. Identities are typically assigned by the system operators or by trusted third parties and used in a centralized or decentralized manner. However, some access control models do not use strict identification of users.

DLT systems can achieve access control goals using a variety of approaches. Services supporting access control can be deployed in the user and non-DLT systems layers, in the development functions, and in the management and operation functions. As an example, DLT systems can implement authorization with or without identification by approving anonymous transactions that meet specified criteria. Mechanisms for this can include tokenization, which uses cryptography to assign rights and privileges to controllers of certain cryptographic keys. The requirement to identify agents for access control purposes depends on the legal or business context.

6.5 Storing records

From a records management point of view, requirements for storing records have changed a lot in the transition from the paper-based environment to the digital one. In the paper environment, storage is directly related with custody or physical possession of a record, but in the digital environment storage became more a technical question related to servers, networks or storage devices. The scope of ISO 15489-1 covers both environments and describes only a limited set of measures to ensure appropriate storage environments and media:

— the use of protective materials and special handling procedures where necessary;

— routine protection and monitoring of physical and information security;

— development and testing of authorized disaster planning and recovery procedures;

— the training of relevant personnel in these measures.

Routinely monitoring and evaluating these measures to identify any risks to the records' accessibility or integrity, are the actions related to the process of storing records. ISO 15489-1 doesn't provide advice for the case when the records are stored in multiple copies and there is no designated official record.

There are many different records storage configurations that can exist in the context of DLT systems. Records can be:

— stored on-ledger;

— stored off-ledger, and linked to metadata stored on-ledger (see 5.4);

— stored off-ledger, in a centralized data store, e.g. cloud based database;

— stored off-ledger in a decentralized data store, e.g. InterPlanetary File System [IPFS].

In the design of DLT systems, an important decision about storage is whether records and associated metadata are created and stored on-ledger or off-ledger. In addition, as for any other system, the designer can make a different choice about infrastructural components of storage (e.g. IPFS, distributed databases or cloud storage). Given the decentralized architecture of DLT systems, records can be scattered across a broad array of systems and infrastructural components. Some of these can be under the control of a single organization, others under the control of business partners who are members of the consortium, and still others under control of unknown third-party agents.

In DLT systems, any data stored on-ledger is replicated on at least some of the DLT nodes that participate in the operation of the network. The decentralized characteristics of a DLT system can be considered an advantage when compared with centralized systems because copies of the ledger records are stored in different DLT nodes. This can increase the likelihood that the data would remain accessible but does not necessarily ensure their authoritativeness nor their long-term preservation. In addition, assessing the authoritativeness of ledger records can be problematic if a DLT system is no longer fully operational. Long term accessibility of authoritative records involves considering in the design phase actions such as contingency plans, metadata on technical dependencies, and monitoring of storage.

DLT systems face some risks similar to those in cloud computing environments. ISO/TR 22428-1 describes a potential risk as follows: “The integrity of a record requires stable storage for as long as the record is required to be maintained. The configuration of virtual servers used in cloud computing changes frequently in order to provide elasticity of service provision on demand. These frequent changes can result in unintended consequences such as unexpected alterations to configuration of record stores, metadata or security controls.”

Preservation of authoritative records in the long term can require extensive data storage. For this reason, records systems sometimes provide functionalities to move semi-active or inactive digital records to less expensive storage media or to off-site storage. In the context of DLT systems such practice can undermine the integrity of the distributed ledger. To reduce storage space, DLT systems sometimes use different storage compression techniques (e.g. pruning of the blockchain). These techniques can present challenges for accessibility of transaction records or the ability to check records integrity.

Apart from the technical issues, storing records in DLT systems is also related to legal issues concerning custody and ownership. These legal issues are discussed in 9.2.3.

6.6 Use and reuse

The useability of records for as long as they are retained is a crucial issue from a records management point of view. As explained in 5.5; a useable record can be located, retrieved, presented and interpreted over time.

Technical dependencies can be a problem for records with a long period of retention. Records management practices include potential actions to mitigate the risk of inaccessibility such as applying and maintaining appropriate metadata about a record’s technical dependencies, creating additional copies of records, or converting them into alternative formats, migrating records, preparing a plan to ensure continued access and useability of records in the event of a disaster affecting records systems or storage areas, or establishing routine monitoring of storage conditions.

Use of the ledger records outside of the DLT system presents several challenges relating to proving the characteristics of authoritative records. For example, when needing to extract records from a ledger for purposes of presentation to governments, courts or third parties, it can be difficult to preserve the relationships among ledger records and between ledger records and other records stored off-ledger. The inability to establish relationships among records limits a user’s ability to demonstrate a record’s authoritative characteristics.

Possible solutions to establish or guarantee the authenticity of ledger records made available for use outside of their originating DLT system can include but are not limited to:

— explicitly setting the procedure for obtaining an authenticated record copy of the ledger records in legislation or in contracts and agreements. Procedures for authenticating record copies are likely to vary by jurisdiction according to applicable legislation, standards and custom but, for example, can include re-hashing and re-signing/re-sealing previously digitally signed records, placing records in a trusted third

party digital repository (see ISO 17068) for subsequent use, or enacting legal provisions that recognize a ledger record registered in a DLT system as authentic, provided that it is accompanied by a written declaration of a qualified person, made under oath;

— allowing the relying parties or authorities (e.g. regulators, courts, public notaries etc.) direct access to a DLT system currently in use;

— interoperability with authoritative external systems (e.g. government registers) with requisite security measures.

The presentation of ledger records to authorities or courts in legal proceedings is further discussed in 9.2.

6.7 Migrating and converting records

Migration and conversion of records are the most common practices related to digital preservation of long-term retained records. Migration is about relocating records from one system to another for different reasons such as the decommissioning of the original system. Conversion is the change of formats due to the obsolescence of the original ones. Guidelines of how to perform these records processes maintaining the characteristics of authenticity and reliability are provided in ISO 13008.

Given that operating a multitude of deprecated DLT protocols is likely to place a strain on any system, there is a need of a migration strategy and planning to transfer a record created and recorded in one DLT system into a record on another system without losing the authenticity and reliability.

Conversion of the format of a record results in the change in its hash value within the DLT systems. The tamper resistance of DLT technologies relies on the security of the hashing algorithm in both on-ledger and off-ledger systems, so in principle any conversion of formats is a complex challenge.

Since distributed ledgers are designed to be immutable, conversion of records, whether stored on or off-ledger, can be problematic. Approaches to conversion in the context of DLT systems are discussed in 9.9.

6.8 Disposition

According to ISO 15489-1:2016, 9.9, “Records and metadata should be retained for the time periods specified in disposition authorities” and “disposition processes should be carried out in conformance with rules in authorized and current disposition authorities.”

Records systems are expected to support disposition processes and actions by providing corresponding functionality. DLT systems are typically designed to create final, definitive and immutable ledger records for which disposition actions are not contemplated. In such cases, DLT systems do not support disposition processes and actions even when there is a legal or operational requirement for records disposition (see 9.1). There can be a risk of not being able to delete all copies of the information in a DLT system.

In other cases, where DLT systems are designed to support disposition actions and processes using a variety of techniques (e.g. cryptographic shredding, chameleon chains, etc.), this capability has the potential to weaken the authoritativeness of the ledger. DLT system designers would typically consider the trade-offs in implementing disposition functionality.

In certain cases, such as transfer of records to an archival institution for permanent preservation, disposition can involve transfer of the control of records and metadata to another system user. Well-defined processes for execution of such disposition actions have yet to emerge in relation to DLT systems.

ISO 15489-1:2016, 9.9 specifies:

“The following principles should govern the destruction of records:

a) destruction should always be authorized;

b) records pertaining to pending or actual litigation or legal action or investigation should not be destroyed while that action is underway or anticipated to arise;

c) records destruction should be carried out in a way that ensures complete destruction and which complies with any security or access restrictions on the record;

d) destruction, like any disposition action, should be documented."

Regarding the above recommendations, DLT systems, which often support process transparency, can function very effectively to track authorizations for any off-ledger destruction actions, protect records relating to pending litigation or legal action from destruction, and documentation of destruction actions. Further, smart contract capabilities of some DLTs can also be used to automate off-ledger destruction processes and actions.

7 Relationship between distributed ledger technology (DLT) systems and records systems

7.1 Characteristics of records systems

The authoritativeness of records is supported if records are managed in systems that are reliable, secure, compliant, comprehensive, and systematic. This subclause describes the characteristics of records systems, the relationship between the characteristics of records systems and those of DLT systems and discusses the possibility of DLT systems being utilized as records systems.

According to ISO 15489-1 there are five characteristics of records systems.

— Reliable: Records systems should be capable of continuous and regular operation in accordance with authorized policy and procedures.

— Secure: Measures such as access control, monitoring, agent validation and authorized destruction should be implemented to prevent unauthorized access, alteration, concealment or destruction of records.

— Compliant: Records systems should be managed in compliance with requirements arising from business, community or societal expectations and the legal and regulatory environment.

— Comprehensive: Records systems should be capable of managing all required records of the range of business activities to which they relate.

— Systematic: The creation, capture and management of records should be systematized through the design and routine operation of records systems, and by adherence to authorized policies and procedures.

Records systems provide measures to prevent unauthorized access, alteration, concealment, or disposition of records, and to enable the participation of any authorized agents. Some DLT systems have limited capabilities for authentication and access control by design, so additional system elements can be needed to use such DLT systems as records systems. Some public or permissionless DLT systems provide limited identity and access management that can affect the desired records systems characteristics of reliability and security, with implications for the authoritativeness of records. In the case of permissioned DLT systems, it can be possible to restrict the participation in the DLT network. However, unlike records systems in which various levels of access to records are assigned to each entity, DLT nodes in some DLT systems can have the same access right to records on distributed ledgers which can have privacy and security implications.

For reliability, records systems provide mechanisms, when necessary, for importing (or otherwise incorporating) records and metadata for records into the system or exporting them from one system to another. However, as with other technologies it is a challenge to export records and metadata for records from a DLT system to another since there is no universally recognized exporting mechanisms that preserves the authoritativeness of records. As DLT systems are designed to ensure immutability of the records contained in a distributed ledger, then in case of export, it can be necessary to establish an explicit procedure to determine which instance of records are authoritative.

Records systems facilitate disposition actions including destruction or transfer of records and their associated metadata. Records disposition sometimes involves the complete destruction of records in compliance with legal and other relevant requirements. However, distributed ledgers are designed to be immutable, so that ledger records would not be able to be altered or deleted. Furthermore, in circumstances

wherein there is a legal requirement to delete a record, it is sometimes more difficult or impossible for authorities to enforce that requirement for all participants in a decentralized system. If a records system stores records as ledger records in a DLT system, and if the DLT nodes are operated by parties who are not likely to comply with a legal requirement to delete those records, then it will be more difficult to satisfy that legal requirement.

The inability of a DLT system to destroy records can conflict with legal requirements that specify the right to correct or delete certain information such as personal information, offensive information, or other information required by courts.

Some DLT systems have records that are distributed across different jurisdictions, and this can conflict with legal requirements for data localization that data must be stored in a specific jurisdiction. This can affect the compliance of records systems with relevant requirements.

7.2 Design considerations for records systems

ISO/TS 16175-2 provides guidance for decision making and processes associated with the selection, design, implementation and maintenance of software for managing records, according to the principles in ISO 15489-1. The following list maps the principles from ISO 15489-1 to the clauses in this document to provide context for the design of DLT systems.

— The creation, capture and management of records are integral to conducting business, in any context (for details in the context of DLT, see Clause 6).

— Records, regardless of form or structure, are authoritative evidence of business when they possess the characteristics of authenticity, reliability, integrity and useability (for details regarding DLT and authoritative records, see 5.5).

— Records consist of content and metadata, which describes the context, content and structure of the records, as well as their management through time (for details in the context of DLT, see 5.4 and 8.2.2).

— Decisions regarding the creation, capture and management of records are based on the analysis and risk assessment of business activities, in their business, legal, regulatory and societal contexts (for details in the context of DLT, see 4.2, 9.1 or 9.3).

— Systems for managing records, regardless of their degree of automation, enable the application of records controls and the execution of processes for creating, capturing and managing records. They depend on defined policies, responsibilities, monitoring and evaluation, and training in order to meet identified records requirements (for details in the context of DLT, see 7.1, 8 and 9.1).

The application of these principles varies by the kind of DLT system and its utilization. For details concerning interaction of DLT-systems and records systems in records management, see 7.1 and Clause 8.

The DLT system itself is used to manage on-ledger records. Off-ledger records can be managed by a non-DLT records system as well. At the same time, DLT systems can play a role in the management of off-ledger records by supporting such processes as protecting the integrity of records, access control, retention, and so on. In some DLT systems, it can be more difficult to manage off-ledger records than in conventional records systems. External records systems can be used in combination with DLT systems to help to address those challenges.

The design and implementation of records systems based on DLT will benefit from the comprehensive analysis of the context of records creation and management (e.g. the legal, business, technical and other applicable context). In some jurisdictions, specific measures can be necessary to satisfy the legal and regulatory requirements for on-ledger records.

8 Distributed ledger technology (DLT) systems and records management

8.1 Policies and responsibilities

The purpose of records policies is to support requirements for the creation, capture, and management of authoritative records. Policies can also exist for the design, use, and management of records systems. In the DLT context, policies set the high-level framework for DLT system's functioning and define the roles and responsibilities for records.

Clear and well-defined roles and responsibilities will support records creators, those ones involved in management of the records and other users of records systems. Appropriate monitoring and evaluation ensure for example the identification for changes, further training and so the continuous improvement of this records management governance.

Depending on the DLT system, the specific manifestation of policies and responsibilities differs. In private permissioned DLT systems explicit rules concerning the creation, management, storage, and use of records can be specified in policy documents and implemented by the operators or operator consortium. Content of the policies includes scope, relevant and applicable law and standards, possible auditing requirements, related business activities, and procedures. Approval and promulgation can be done by a designated person in the operator's side of permissioned DLT systems. Well-defined and assigned responsibilities ensure establishment and improvement of proper records management. This can set out the responsibilities for specific roles for development and implementation of records policies, guaranteeing that records are managed to meet business needs, and ensuring appropriate operation, security, performance, and scalability of DLT as records systems. All DLT users bear some responsibility for creation and maintenance of the shared distributed ledger.

In some permissionless DLT systems, system-wide records policies are only promulgated, implemented, and enforced by dedicated technical measures, such as by means of software or algorithms (e.g. Nakamoto consensus mechanisms).

8.2 Records controls

8.2.1 General

Records controls assist in meeting records requirements and include metadata schemas for records, business classification schemes, access and permission rules, and disposition authorities. Records controls can be designed and implemented in a variety of forms, depending on the technological and business environment. In defining and implementing records controls in the context of a DLT system, relevant considerations include the design of the DLT system, DLT system operating requirements, and legal and regulatory requirements of relevant jurisdictions of system operation.

8.2.2 Metadata schemas

Metadata schemas are developed to define the metadata used to identify, describe and manage records. In the context of DLT systems, some metadata can form part of the contents of a ledger record. For ledger records to possess the characteristics of authoritative records associated metadata ideally is based on a standard metadata schema. Because DLT systems are decentralized, reaching agreement on a standard metadata schema or on subsequent changes to an existing schema can require considerable coordination and effort. In some jurisdictions, pre-existing, authorized metadata schemas for a jurisdiction can be used or adapted assuming that the DLT system operates in only that jurisdiction.

Metadata schemas can relate to different entities. Key entities for managing records are the following:

— Records – including all levels of aggregation;

— Agents – including persons, business units, technologies or business and records systems;

— Business (or Function) – business functions, activities and transactions or work processes;

— Mandates – laws and other requirements governing the conduct of business and record creation or management;

— Relationships – between entities and layers of aggregation.

According to ISO 15489-1:2016, 8.2:

“Metadata should be defined to:

a) enable the identification and retrieval of records;

b) associate records with changing business rules, policies and mandates;

c) associate records with agents, and their authorizations and rights with regard to the records;

d) associate records with business activities;

e) track processes carried out on records, such as changing access rules or migrating records into new systems.”

Six broad classes of metadata are typically used in the management of records. They can be applied to all entities (see above), or some, depending on the complexity of the implementation. The six classes are the following, several of which are not commonly found in DLT systems at the present time.

Table 1 — Metadata classes and distributed ledger technology (DLT) context examples

Metadata classes	Explanation from ISO 15489-1	DLT context examples
Identity	Information to identify an entity	A DLT address for identification of an agent or a transaction hash for identification of a record
Description	Information to determine the nature of the entity	Metadata referencing an off-ledger ontology or business classification scheme
Use	Information that facilitates immediate and longer-term use of the entity	Not typically included in DLT systems
Event plan	Information used to manage the entity, such as disposition information	Not typically found in DLT systems, likely owing to the association of the concept of immutability with permanent retention
Event history	Information recording past events on both the entity and its metadata	Not typically found in DLT systems, because any changes to metadata found in ledger records constitute an alteration that will render those records invalid (i.e. lacking integrity)
Relation	Information describing the relationship between the entity and other entities	Metadata referencing an off-ledger ontology or business classification scheme

8.2.3 Business classification schemes

Business classification schemes are tools for linking records to the context of their creation. By linking records requirements to a business classification scheme, processes for the appropriate management of records can be carried out.

The act of linking a record to its business context is the process of classification, which supports the following:

— the application of access and permissions rules;

— the execution of appropriate disposition rules;

— the migration of records of a particular business function or activity to a new environment as a result of organizational restructure.

In the context of DLT systems, it can be desirable from a records management point of view to link the DLT records with relevant business classification scheme(s). Links can be created using several different approaches, such as via embedding metadata representing an ontology or business classification scheme identification and version number into a ledger record.

Like development of metadata schemas for DLT systems, reaching agreement on a standard business classification scheme or on subsequent changes to a scheme can require considerable coordination and effort. In some jurisdictions, pre-existing, authorized business classification schemes for the jurisdiction can be adhered to or adapted assuming that the DLT system operates in only that jurisdiction.

8.2.4 Access and permissions rules

A set of rules identifying rights of access and the regime of permissions and restrictions applicable to records typically exist for all records systems.

Access and permissions rules usually can be associated with the identification of agents, and their rights and permissions. In permissionless DLT systems, however, there exist no agent-specific rights or permissions. In contrast, such permissions can be defined in permissioned DLT systems.

In both permissionless and permissioned DLT systems, rights and permissions concerning business activities or access to records can be defined. For example, these rights and permissions can include who can view specific components of ledger records or specific types of transactions that are only viewable by specific agents.

8.2.5 Disposition authorities

Disposition authorities authorize the disposition of records and are required for conformity with ISO 30301 or necessary for legal or business reasons. In the context of DLT systems, which are designed to create final, definitive and immutable ledger records, the implementation of disposition authorities can appear to be in conflict with the design goals of DLT systems. In such cases, the design and implementation of disposition authorities would address records requirements including authoritative records, the design and governance of the DLT system, and the legal and regulatory requirements of the relevant jurisdictions.

Disposition authorities typically are authorized, dated, implemented and regularly reviewed to take account of changing requirements. In some jurisdictions, this often involves authorization from an external authority or regulator. Authorization by a sole external authority or regulator is likely to be inadequate in DLT systems operating across multiple jurisdictions.

Implementation of disposition authorities ideally is monitored and documented, and regularly reviewed for any new requirements affecting the business activity documented in the records. In DLT systems, monitoring of the implementation of disposition authorities can prove challenging as access to decentralized instances of the ledger cannot be guaranteed. Compliance with disposition authorities can also be uncertain, as it is always possible for DLT users to retain off-line copies of ledger records without the knowledge of other DLT users or DLT system operators.

The disposition authorities themselves can be managed in DLT-based systems, regardless of whether the records are managed in that DLT system.

9 Challenges, considerations and potential benefits

9.1 Distributed ledger technology (DLT) and management of retention and disposition of records

ISO records management standards provide a model for retention and disposition of records based on the appraisal of business context to determine which records need to be created and captured and how long the records need to be retained. The appraisal process can also be used for determining how to apply records processes to the records. Appraisal is context-dependent, and as a result, not every record is managed in the same way (see ISO/TR 21946).

The appraisal process can result in a definition of records requirements and disposition process, and consequently disposition authorities (see 6.8 and 8.2.5), that can be implemented in DLT systems, when the characteristics and constraints for ledger records are understood through the analysis of the business and technological context.

Depending on the specific context of the records existence in DLT systems, it can be challenging to apply the concepts of records creation, capture and retention. For example, creation and capture processes typically are not as distinct in DLT systems as they are in other types of records systems. In DLT systems, new ledger records are not independent, because they are cryptographically linked to pre-existing ledger records.

Applying the concepts of records retention and disposition involves considering the properties of DLT systems. In DLT systems, ledger records are meant to be final, definitive and immutable (i.e. permanent) once they are recorded on the distributed ledgers in those systems. It is usually not possible to destroy individual ledger records without compromising the integrity of the distributed ledger and the authoritativeness of the remaining ledger records.

It is possible to transfer the ledger in its entirety from a DLT system to another system without affecting the integrity of its ledger records. The partial transfer of ledger records from a DLT system to another system is possible, however, this is likely to impair the authoritativeness of the records. Tagging ledger records designated for permanent retention or destruction is one way of preparing the ledger for subsequent disposition.

9.2 Legal issues

9.2.1 General

DLT systems can present several legal issues. One of the main legal issues facing the DLT systems is that of jurisdiction. In some cases, DLT systems can operate across jurisdictions. Due to their architecture and manner of operation, it can be difficult to determine which jurisdiction is applicable in relation to resolving disputes or enforcing contracts.

Compliance with various laws and regulations can also prove challenging. For instance, in DLT systems that do not require identities for execution of transactions (e.g. permissionless DLT systems), compliance with anti-money laundering (AML) and know-your-customer (KYC) rules can be more difficult. However, use of decentralized identifiers and verifiable credentials associated with Self-Sovereign Identity (SSI) can sometimes be used to address this issue without reliance on the types of access control mechanism typical of centralized systems. Compliance with data privacy and data protection laws can also be an issue, depending on the design of the DLT system. In a similar manner, DLT systems can be used to facilitate cross-border transactions that bypass existing regulatory controls. Depending on its design, it is possible that a DLT system does not rely on traditional measures (e.g. associating rights with specific identities) for the prevention of those transactions. Smart contracts can be used to mitigate the risk of occurrence of problematic transactions (e.g. by using explicit programmed checks to verify that the transactions are allowed). Once a transaction has been added to a distributed ledger, it cannot generally be deleted or modified. However, it is possible to create an amending transaction to reverse an earlier transaction.

Uncertainty about intellectual property rights can also present a legal issue. Furthermore, in some jurisdictions, a compilation of data (i.e. a distributed ledger) falls within the scope of intellectual property rights, and thus can require participants in a DLT system to clarify the IP rights associated with ledger records in a DLT system.

An additional legal issue concerns liability, especially in public permissionless systems. As an example, since smart contracts execute across many nodes in a DLT system, it can be difficult to determine who is responsible in the event of an error or breach. For example, if a smart contract is hacked and funds are stolen, it can be unclear who is liable for the loss. Relatedly, if a bug is discovered in smart contract code, it can be difficult to identify and agree who is responsible for fixing it when there is no clear central authority responsible for the operation of the associated DLT system. Legal disputes are not likely to be resolved by traditional mechanisms. New approaches, such as decentralized arbitration services or dispute resolution mechanisms embedded in smart contracts can be used to help to address this challenge. In addition, when problematic transactions have occurred, it can be possible to reidentify pseudonymous parties using analytic techniques to hold them accountable.

9.2.2 eDiscovery

Where electronic discovery, or eDiscovery, exists, the process typically involves identification of records pertinent to the legal matter, and preservation, collection, processing, review, analysis, production, and presentation of those records as evidence. DLT systems can present unique challenges for eDiscovery. In terms of identification of relevant records, the decentralized nature of DLT systems can make it challenging to identify and collect authoritative records. Unlike traditional centralized systems, where it can be possible to obtain records from a single centralized source (e.g. a central registry), ledger records exist on numerous DLT nodes, which are not necessarily identical and often exist across jurisdictions, which can require coordination and cooperation from multiple parties to access the relevant evidence. In addition, DLT networks can be operated without a central authority or administrator. This lack of central control can complicate the process of identifying custodians of the relevant ledger records and determining who can provide access to the evidence. Unlike traditional systems where organizations or individuals have control over their records, DLT network participants can have limited control over the ledger records they generate, making it difficult to establish responsibility and accountability. Some of these issues can be addressed by paid services that use smart contracts to identify and collect copies of relevant ledger records. Understanding and interpreting records from DLT systems involves specialized knowledge and expertise. The cryptographic algorithms, data structures, and consensus mechanisms involved in DLT can pose a challenge for those who are not familiar with these concepts. Bridging the gap between legal and technical domains becomes essential for successful eDiscovery in DLT-related cases.

In terms of preservation, collection, processing, review, analysis, production, and presentation of records as evidence, DLT system transactions often use cryptographic addresses instead of personal identifiable information (PII) associated with identities. This pseudonymity can make it challenging to link ledger records to real-world identities during eDiscovery. Connecting DLT system addresses to individuals or organizations involves additional investigation, and without cooperation from involved parties, it can be challenging to establish a clear chain of custody for the evidence. On the other hand, DLT's immutability, a key feature that ensures the integrity of the data stored on-ledger, means that it becomes nearly impossible to modify or delete evidence, thereby offering greater capacity for protection and preservation of evidence in DLT systems. The sheer volume of data in many DLT systems can present challenges for analysis in eDiscovery. Analysis tools can help alleviate this burden in DLT-based records systems. The legal admissibility of ledger records in relevant jurisdictions can present challenges. Given the distributed nature of DLT systems, identifying the location of the authoritative copy of a ledger record can be difficult. Submitting DLT ledger records into evidence therefore can be challenging because of the existence of multiple and potentially non-identical copies (which can call into question a claim of immutability) and lack of a single “official” record copy (which is a requirement in some jurisdictions). This issue can be especially difficult when government bodies or courts are involved. This raises the issue of selecting the appropriate method of submitting DLT-based records as evidence to courts, as this is jurisdiction dependent.

9.2.3 Custody and ownership

Traditionally, records have been most often managed in the framework of a single organization that owned and controlled them, even if the storage and custody of those records has been outsourced to a third party. In the case of DLT systems, records can be created and stored in a distributed manner in the absence of centralized control. Consequently, records created or received by a DLT system can be in the custody of, and under the control of, participating DLT nodes that are owned and operated by different, independent participants in a DLT network. Participants can be identifiable or pseudonymous organizations or individuals, depending on the type and design of the DLT system. Unlike traditional information systems, some DLT systems can have no designated owner or controller, resulting in legal uncertainties about ownership, custody and legal responsibilities. DLT node owners and operators can potentially each claim ownership rights to the records generated by and held within a DLT system. This can create some legal ambiguity regarding the ownership of such records as well as concerning participants’ legal responsibility for compliance with laws governing the handling of such records (e.g. laws on data protection, copyright, and preservation of evidence). Finally, identification of ownership rights can be difficult to establish in DLT systems wherein participants operate DLT nodes pseudonymously. Consequently, participants in a DLT system can require clarification of the custody and ownership of records generated, received by, and stored in DLT systems.

Additional custody and ownership issues can arise when a DLT system is operated across several jurisdictions. These issues are discussed in 9.3.

9.2.4 Geolocation restrictions on data storage and transfer

Knowledge of relevant laws and regulations is essential for understanding restrictions on, for example, records content, formats, and storage location.

There are considerations for DLT systems and their operation across jurisdictions. Laws and regulations can specify requirements for the location of systems that store or process certain data, or for the transfer of data across borders.

Knowledge of relevant laws and regulations is also essential when the operation of a DLT system results in data being shared or transferred across jurisdictional boundaries.

There also can be implicit geographical location requirements for records, such as those found in policies requiring that records must be in the principal place of business or employment.

Compliance with geolocation-related restrictions on data location and transfer can be challenging due to the distributed nature of data transmission and storage in some DLT systems.

9.2.5 Jurisdictional restrictions on the operation of distributed ledger technology (DLTs)

Operation of DLT systems can be subject to laws and regulations of applicable jurisdictions. In some jurisdictions, only dedicated service providers are allowed to operate or provide a DLT system for certain application areas, such as national identity management, mining, banking, and so on. There can be jurisdictional technical requirements and guidance related to the use of cryptography in the DLT systems. Depending on the application area of the authoritative records (e.g. privacy, data exchange, governance, etc.), geographical restrictions can apply. If a DLT system has DLT nodes that operate in different jurisdictions, or is used for cross-border transactions, then the laws and regulations of the relevant jurisdictions can apply.

Conducting a well-grounded requirements analysis before using DLT systems as part of records systems or as records systems themselves can address these challenges.

9.3 Personally identifiable information (PII) protection

High-level principles of privacy and PII protection and corresponding challenges are described in ISO 23257 and ISO/TS 23635. According to ISO/TR 23244, the principles are:

— consent and choice;

— purpose legitimacy and specification;

— collection limitation;

— data minimization;

— use, retention and disclosure limitation;

— accuracy and quality;

— openness, transparency and notice;

— individual participation and notice;

— accountability;

— information security;

— privacy compliance.

PII protection requirements have a significant impact on records processes, especially on access control, retention and disposition of records. The severity of PII protection issues differ depending on the legislation of applicable jurisdictions and on business context. PII legislation can sometimes override other legislation concerned with the management or preservation of records.

PII protection issues are widely seen as a major barrier for the adoption of DLT-based solutions. The major challenge is ensuring compliance with PII expungement requirements of relevant privacy and personal data protection legislation, given that distributed ledgers are immutable by design. As stated in ISO/TR 23244, "…Modifying, deleting or adding information or transactions can be difficult on a blockchain or DLT system as this can destroy the integrity and immutability of the ledger; also, it can be difficult to gain agreement between users, operators and administrators to modify, alter or add to the ledger; and finally, the system may not have the capabilities to perform such activities."

However, there can be other specific DLT-related issues that are legal in nature.

— Identification of relevant jurisdictions and applicable laws. As explained in ISO/TR 23244, "Blockchain and DLT systems can involve many stakeholders living and working in different countries and different legal and regulatory environments. The challenge for a blockchain and DLT system and its stakeholders is to provide legal certainty through enforceable agreements, contracts and associated mechanisms, under an agreed and recognized legal jurisdiction." As a result, identifying legislation applicable to the operation of a given DLT system can be difficult.

— Identification of the responsible parties. As also stated in ISO/TR 23244, "A further challenge is that as some blockchain and DLT systems could not have a clearly defined 'owner' or be a clearly identified legal entity, it can be difficult to apply the accountability principle as laid out in ISO/IEC 29100 and some jurisdictions can have difficulty in interacting with a system without clearly defined legal status… Courts and authorities can require disclosure, deletion, modification or addition of certain information or transactions. Complying with such legal requirements can be difficult for blockchain and DLT systems and their users, operators and administrators. A disclosure request and the disclosed data can identify a PII principal or provide relevant search attributes which can result in non-PII becoming PII, or allow a PII principal to be indirectly identified…. Thus, identification of the operator or owner of a DLT system might not be possible."

— Managing consent for PII processing. PII laws usually provide for legal processing of PII based on consent, as well as several exceptions when consent is not required. Consent is usually easy to revoke. At the same time the legal practice concerning PII processing based on exceptions, is neither uniform across the jurisdictions nor stable in time. For some types of DLT systems that include PII in the ledger, the management of the consent of the PII principals can be difficult to implement.

ISO's records management standards provide little general guidance on PII protection issues. ISO/TR 23244 provides a detailed overview of the issues and practical concerns related to privacy and personally identifiable information (PII) protection in the context of blockchain and distributed ledger technologies (DLT) and their applications.

PII protection issues are sometimes easier to resolve when PII is processed in government-owned or accredited information systems based on the special laws covering these specific systems. The same would be true for government-owned or government-endorsed DLT solutions.

In the case of PII processing in a private permissioned DLT system, there can be exogenous mechanisms, e.g. courts of law, to compel the behaviour of participants to implement PII expungement.

The most challenging PII-related problems are associated with public permissionless DLT systems. Generally, it is considered risky to place PII on ledger in these systems.

When designing a DLT system that is intended for PII processing, it is advisable to consider privacy-by-design and privacy-by-default approaches. Specific guidance on privacy-by-design for DLT-based solutions can be found in DIN SPEC 4997.

The concept of privacy-by-design is to consider data protection from the beginning of the development of an application that processes personal data. The potential risks for the rights and freedoms of natural

persons are typically assessed at early stages and can be mitigated by suitable technical and organizational measures.

The concept of privacy-by-default is to implement privacy-friendly settings by default, instead of settings that allow extraction of more personal data than needed. In this regard, the public DLT systems that make personal data transparent for everyone pose a problem to be mitigated.

It is also recommended to conduct, where appropriate, a privacy impact assessment (PIA), which is the overall process of identifying, analyzing, evaluating, consulting, communicating and planning the treatment of potential privacy impacts with regard to the processing of personally identifiable information, framed within broader risk management framework. Guidance on PIA can be found in ISO/TR 23244, ISO/IEC 29134 and ISO 22307.

Typical privacy threats to be considered include the following.

— Uncontrolled access to information and PII.

— Accidental or deliberate exposure of PII.

— Poor implementation of security technologies, including cryptography.

— Loss or publication of cryptographic keys.

— Loss or publication of access credentials.

— Exploitation of obsolete or out-of-date hardware, middleware and software.

— Attacker writing sensitive PII into the ledger.

Typical vulnerabilities of the DLT systems include the following.

— Poor password management (to include using default passwords).

— Lack of access management.

— Poor patching and updating processes.

— Poor coding practices (to include the use of backdoors).

— Poor user training.

— Poor physical security.

In many cases, privacy information management systems (PIMS) can address the management of personal information that is held across a wide range of operational units and information technology-based application systems. ISO/IEC 27701 offers guidance for governance models for PIMS. The key issue from a DLT systems' perspective is how to put in place a PIMS implementation across multiple organizations and jurisdictions.

9.4 Access control mechanisms

9.4.1 General

Records systems are implemented to, among other things, satisfy business or legal needs to control access to records. This is done by controlling write access to records for purposes of controlling records creation and disposition, and to protect privacy and confidentiality by controlling read access for the discovery, viewing, and analysis of records.

All DLT systems have capabilities to control write access to data to support data integrity for ledger records. However, it can be more complex for some DLT systems to control read access to data to support privacy or confidentiality of ledger records. A typical mechanism used in DLT systems to support data integrity is to enable many separate parties to each cross-check the validity of ledger records. Often this is achieved by

those parties having unfettered read access to ledger records and being able to separately retain their own full copy of the ledger.

9.4.2 Read access

In some DLT systems, for example, public permissionless DLT systems, any member of the public can obtain and read a copy of the entire ledger. Some DLT systems are private or permissioned, wherein read access is limited to authorized entities. Both are broad categories, and whether a specific private or permissioned DLT system is adequate for supporting privacy or confidentiality requirements for a particular solution will depend on the business needs and design of that solution. Private DLT systems are only accessible to a limited group of users. For example, a private DLT system can operate only within the context of a company's virtual private network which is only accessible to the company's employees. However, it is possible that even a private DLT system does not achieve sufficient read access control. Permissioned systems require authorization to perform certain activities in the system. While those authorized activities can in principle include reading ledger records, specific mechanisms (e.g. private channels) will normally be used to support read access control. DLT systems can be used as components in solutions that have privacy or confidentiality requirements. Those solutions are normally designed with specific additional mechanisms used for read access control. Some of these mechanisms, and their limitations, are identified below.

The most straightforward mechanism for read access control is to avoid recording PII or confidential data onto a distributed ledger, but instead to keep those data off-ledger and control read access to that data set using conventional technologies. This technique can result in loss of data integrity protection provided by DLT systems. Some data integrity support can be recovered by instead storing hashes, which are calculated using cryptographic hash functions on the off-ledger data. Further discussion of security and confidentiality in DLT systems is found in other standards documents, including ISO 23257 and ISO/TR 23244. Further discussion of the security of cryptographic hash functions is found in ISO/IEC 10118-3.

Another technical mechanism for read access control in DLT systems is the use of encryption. Data can be encrypted before being stored on-ledger, and if using a sufficiently strong encryption scheme, the data will remain confidential as long as the decryption key remains secret to the parties who are authorized to read that data. Encryption can be used to implement blind signatures (see ISO/IEC 18370-1), which in addition to supporting confidentiality, can also support integrity for some parties. Some sophisticated encryption schemes, such as homomorphic encryption (see ISO/IEC 18033-6), support confidentiality while also allowing some integrity properties to be independently cross-checked by all parties who can access the distributed ledger. However, these sophisticated encryption schemes tend to be more complex to implement and administer and can impact the performance efficiency of the solution.

Some DLT systems allow transactions to be made by pseudonymous parties. For example, DLT users are able to create their own unique keys in a public key cryptography scheme and use their public keys as a basis for their user identifier in transactions. This can help to support privacy or confidentiality if the other contents of the transaction data recorded on the ledger do not reveal PII or confidential information.

A potential limitation of all the mechanisms above is possible reidentification (see ISO/IEC 27559), attacks using data and metadata of ledger records. Ledger records are created as the output of transactions that capture metadata that includes at least the time when they were created and the entity that created them. Often there are other data recorded such as counterparty identities, related transactions, and plaintext data and metadata necessary to complete the transaction that can also lead to a reidentification (see ISO/IEC 20889).

9.4.3 Write access

In the operation of a DLT system there are three main points of write access control: the transactions submitted by users to create new ledger records, the ledger updated by DLT node operators based on those submitted transactions, and the DLT system itself. The typical mechanism for write access control for user transactions is that those transactions are cryptographically signed using the user's private key and that the DLT node operators and DLT system check that the signature properly authorizes the submitted transaction. Specific DLT systems will perform additional checks of the validity of submitted transaction data. These validation checks are defined and implemented specifically for each DLT system. DLT node operators use the DLT system to perform these checks before using the submitted transaction data to update the ledger by

adding a new ledger record. Smart contracts can also be used to implement application-specific write access control rules in transactions.

DLT oracles are used to add external data to DLT systems, but this is usually done as part of normal DLT system transactions, and so write access control mechanisms implemented for that system can be used as they would be for other transactions. For example, smart contracts can implement write access control rules for DLT oracles.

9.5 Identification, authentication, and authoritative records

The authenticity of authoritative records according to ISO 15489 implies that the identification of entities creating, sending or receiving a record, as well as their authentication, is needed to establish that they really are who they claim to be. Currently identification of natural or legal entities is typically issued and controlled by centralized authorities and mainly focused on the core identity information mentioned in government identification documents, such as name, surname, address, and birth date.

The disadvantage of identities depending only on centralized authorities is that an identified entity depends for issuance and use of its identity on a centralized authority. This means the identified entity is not sovereign in using its identity. The Self-Sovereign-Identity (SSI) paradigm can solve this issue by establishing decentralized identity management where the identified entity obtains control over its identities including core identity and attributes in its wallet without dependence on a centralized authority by default.

DLT systems can be an enabler of SSI due to their decentralized and distributed nature as well as their inherent properties, such as immutability. In comparison to existing identity management used in records systems, the structure of SSI offers the possibility to submit only the information needed to verify a claim about the identity.

Further discussion of the concept of identity and SSI can be found in ISO/TR 23249.

It is also possible to use existing identification procedures to ensure unique identification of entities using the DLT systems as, or in, combination with a record system.

9.6 Addressing the business need to modify records

Records systems are designed to preserve the integrity of records, and for this reason are designed to prevent the unauthorized deleting, overwriting, or updating of records. Nevertheless, there can be a business need for authorized modification or deletion of records stored in records systems for a variety of reasons, including compliance with legal and regulatory requirements. Typically, records systems have mechanisms to allow for such modification or deletion. However, it can be difficult to implement similar measures within DLT systems because of their immutability by design.

If records are stored as ledger records, one possible approach to deal with the business need to modify the records is to introduce a new distributed ledger transaction that amends the previous transaction by way of entering a new, updated transaction record. One challenge with this approach is to ensure that any agent subsequently relying upon the ledger is accessing the latest version of the ledger record. If a DLT system has a function that makes the relationship between records explicit and easy to find and follow, this approach to addressing the need for modifications can be successful.

If the hash values of records are stored as ledger records, an approach to deal with the business need to modify records is to simply add the hash value of the modified record to the ledger in a new transaction.

In case it is desired that the content of the record requiring modification also be removed from the ledger (e.g. because it contains personally identifiable information or errors), processes applicable to the destruction of ledger records also will apply. These processes are discussed in 9.8.

9.7 Distributed ledger technology (DLT) and records destruction

Records disposition ideally occurs as part of routine disposition processes carried out as authorized by current disposition authorities. Records systems are expected to support the execution of disposition

actions. Typically, records and metadata for records will be retained for the time periods specified in disposition authorities and disposition actions will be carried out accordingly in relation to the following:

— Destruction of records and metadata.

— Transfer of control of records and metadata to an organization that has assumed responsibility for the business activity through restructure, sale, privatization or other business change.

— Transfer of control of records and metadata to an institutional or external archive for permanent retention.

Additionally, management of records in compliance with the laws and regulations in many jurisdictions calls for defensible destruction, meaning that destruction of records is transparent, authorized, and compliant with applicable legislation.

DLT systems have certain advantages as well as disadvantages in respect to records disposition. For example, ISO 15489 states that records destruction should always be authorized. However, in DLT systems, wherein authority can be distributed among human or algorithmic entities, traditional records management approaches to establishing disposition authorities, usually the responsibility of a single organization, can be non-transferrable or not implementable. This implies that ledger records can have different retention requirements and retention periods for different entities, e.g. accessing and using the DLT as a shared ledger. Additionally, disposition requirements can vary in different jurisdictions and apply to the same records. Thus, DLT systems can require a records disposition reconciliation mechanism.

Some jurisdictions have legal and regulatory requirements that specify that records pertaining to pending or actual litigation or legal action or investigation not be destroyed while that action is underway or anticipated to arise. The immutability of the ledger records in a DLT system can help to protect records that must be preserved due to pending or actual litigation. Moreover, smart contract capabilities offered by some DLT systems can support automated, compliant and documented destruction of records stored off-ledger once records retention periods have been met. On the other hand, when records are stored on-ledger, the design of DLT systems, in particular the property of records immutability, can prevent records destruction resulting in non-compliance (e.g. with PII deletion requirements in data protection laws; for more information, refer to 9.4). Designers of DLT systems have trialled a range of strategies to address this potential tension, including:

— Making DLT systems editable; that is, relaxing the requirement that DLT records be immutable.

— Breaking the link between off-ledger records and on-ledger metadata that aids discovery of those records to render off-ledger records inaccessible.

— Cryptographic shredding that entails encrypting the records and destroying the private key to render the records inaccessible.

If deletion is not possible or feasible, it is considered good practice to avoid storing PII or other sensitive information on-ledger.

An additional issue with respect to records destruction can be ascertaining whether all copies of records and metadata for records have been destroyed. This issue exists also in cloud records environments. However, in DLT systems, where many more copies of records or metadata for records can exist and can be distributed across various DLT nodes and off-ledger storage repositories, the issue can be amplified, making it more difficult to ascertain that all records and associated metadata have been duly destroyed than with conventional records systems.

9.8 Longevity of distributed ledger technology (DLT) systems

9.8.1 General

As any other information systems, DLT systems become obsolete or stop providing service over time. If a DLT system is intended to support ledger records of long-term value, then the longevity, sustainability and preservation issues are important to consider at the early stages of system design.

9.8.2 Longevity of cryptographic algorithms

One challenge in the utilization of DLT from a records management perspective is the longevity of the strength of the cryptographic algorithms. In a sense, the breaking of the hash function would not be entirely dissimilar to the problem records systems currently face with the expiration of digital signatures to ensure records' authenticity. Longevity of cryptographic algorithms and their outputs over decade-long retention periods can be ensured, for example, by preservation mechanisms based on re-signing and rehashing the relevant data.

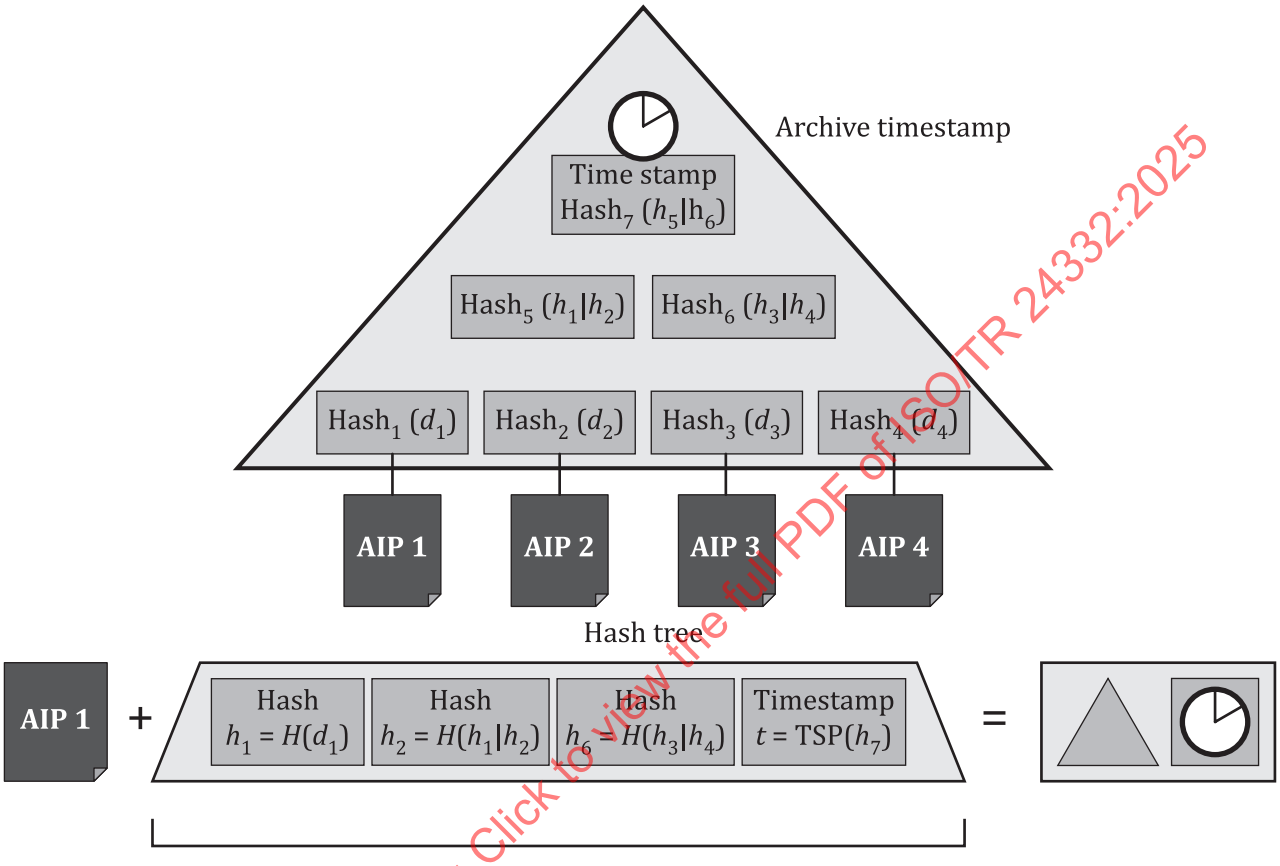

Figure 1 — Hash-tree and evidence record (according to RFC 4998/6283)[1)]

There are potential approaches to applying established rehashing and re-signing procedures to the problem of ensuring longevity of cryptographic algorithms in DLT and international standards have been published in this area.

9.8.3 Long-term preservation of authoritative records

Any approach to long-term digital preservation in a DLT system is complex and warrants careful consideration.

1) The Evidence Record Syntax according to RFC4998/RFC6283 enables processing of several archival information package objects within a single processing pass using a hash tree technique and acquiring only one archive timestamp to protect all archive objects. The leaves of the hash tree are hash values of the data objects in a group. An archive timestamp is required only for the root hash of the hash tree, which ensures efficient processing of large amounts of data. The hash tree can be reduced to a set of hash values, called a reduced hash tree, that is sufficient to prove the existence of a single data object. If the cryptographic algorithms used to create the archive timestamp are at risk of losing their cryptographic suitability, then this archive timestamp can be protected by yet another archive timestamp, which is created with suitable new algorithms before the old algorithms lose their cryptographic strength. For this conservation step, it is necessary to determine whether the relevant signature algorithm or hash algorithm remains secure. Further guidance can be found in ISO 31000 and ISO 27005.