
**Information technology — Open Systems
Interconnection — Distributed Transaction
Processing —**

**Part 1:
OSI TP Model**

iTeh STANDARD PREVIEW

*Technologies de l'information — Interconnexion de systèmes ouverts
(OSI) — Traitement transactionnel réparti —*

Partie 1: Modèle OSI TP

[ISO/IEC 10026-1:1998](https://standards.iso.org/standards/catalog/standards/sist/8d025d6e-b833-4103-a029-2668b178d576/iso-iec-10026-1-1998)

<https://standards.iteh.ai/catalog/standards/sist/8d025d6e-b833-4103-a029-2668b178d576/iso-iec-10026-1-1998>



Contents	Page
Foreword	iv
Introduction	v
1 Scope	1
2 Normative references	1
3 Definitions	2
3.1 Terms defined in other International Standards	2
3.2 Terms defined in ISO/IEC 10026	3
4 Abbreviations	8
5 Conventions	8
6 Requirements	8
6.1 Introduction	8
6.2 User requirements	9
6.3 Modelling requirements	9
6.4 OSI TP Service and Protocol requirements	10
7 Concepts of distributed TP	10
7.1 Transaction	10
7.2 Distributed transaction	10
7.3 Transaction data and coordination level	10
7.4 Tree relationships	11
7.5 Dialogue	11
7.6 Dialogue tree	12
7.7 Transaction branch	12
7.8 Transaction tree	13
7.9 Channel	13
7.10 Handshake	13
7.11 Hinterland	13
8 Model of the OSI TP Service	14
8.1 Nature of the OSI TP Service	14
8.2 Rules on dialogue trees	15
8.3 Rules on transaction trees	16
8.4 Naming	18
8.5 Data transfer	19
8.6 Coordination of resources	19
8.7 Recovery	24
8.8 Concurrency control and deadlock	31
8.9 Security	31

iTech STANDARD PREVIEW
(standards.itech.ai)

<https://standards.itech.ai/catalog/standards/sist/8d025d6e-b833-4103-a029-2668b178d576/iso-iec-10026-1-1998>

Annexes

A Relationship of the OSI TP Model to the Application Layer Structure	Erreur! Si
B Tutorial on concurrency and deadlock control in OSI TP.....	34
C Tutorial on the presumed rollback two-phase commit protocol.....	35
D Combinations of Commitment Optimisations	36
E Summary of changes to the second edition.....	39

Tables

Table 1 - Permitted combinations of transaction data and coordination levels.....	11
Table 2 - Update of log-damage record	24
Table 3 - Types of failures	25
Table 4 - Restoration of node state after atomic action data unavailability.....	30

Figures

Figure 1 - Transaction hinterland of node A viewed from node B.....	14
Figure 2 - Transaction branches, dialogues, and application-associations	18
Figure 3 - Phases of recovery.....	29

iTeh STANDARD PREVIEW
(standards.iteh.ai)

ISO/IEC 10026-1:1998

<https://standards.iteh.ai/catalog/standards/sist/8d025d6e-b833-4103-a029-2668b178d576/iso-iec-10026-1-1998>

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work.

In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75 % of the national bodies casting a vote.

International Standard ISO/IEC 10026-1 was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 21, *Open systems interconnection, data management and open distributed processing*.

This second edition cancels and replaces the first edition (ISO/IEC 10026-1:1992), which has been technically revised. It also incorporates Technical Corrigendum 1:1996.

This part of ISO/IEC 10026 is technically aligned with ITU-T Recommendation X.860, but is not published as identical text.

ISO/IEC 10026 consists of the following parts, under the general title *Information technology — Open Systems Interconnection — Distributed Transaction Processing*:

- Part 1: *OSI TP Model*
- Part 2: *OSI TP Service*
- Part 3: *Protocol specification*
- Part 4: *Protocol Implementation Conformance Statement (PICS) proforma*
- Part 5: *Application context proforma and guidelines when using OSI TP*
- Part 6: *Unstructured Data Transfer*

Annex A forms an integral part of this part of ISO/IEC 10026. Annexes B to E are for information only.

Introduction

ISO/IEC 10026 is one of a set of standards produced to facilitate the interconnection of computer systems. It is related to other International Standards in the set as defined by the Reference Model for Open Systems Interconnection (ISO/IEC 7498-1). The Reference Model subdivides the area of standardization for interconnection into a series of layers of specification, each of manageable size.

The aim of Open Systems Interconnection is to allow, with a minimum of technical agreement outside the interconnection standards, the interconnection of computer systems

- a) from different manufacturers;
- b) under different management;
- c) of different levels of complexity; and,
- d) of different technologies.

ISO/IEC 10026 defines an OSI TP Model, an OSI TP Service and specifies an OSI TP Protocol available within the Application Layer of the OSI Reference Model.

The OSI TP Service is an Application Layer service. It is concerned with information which can be related as distributed transactions, which involve two or more open systems.

ISO/IEC 10026 provides sufficient facilities to support transaction processing, and establishes a framework for coordination across multiple OSI TP resources in separate open systems.

ISO/IEC 10026 does not specify the interface to local resources or access facilities that are provided within the local system. However, future enhancement of the standard may deal with these issues.

iTeh STANDARD PREVIEW
(standards.iteh.ai)

ISO/IEC 10026-1:1998

<https://standards.iteh.ai/catalog/standards/sist/8d025d6e-b833-4103-a029-2668b178d576/iso-iec-10026-1-1998>

Information technology — Open Systems Interconnection — Distributed Transaction Processing —

Part 1: OSI TP Model

1 Scope

This part of ISO/IEC 10026:

- a) provides a general introduction to the concepts and mechanisms defined in ISO/IEC 10026;
- b) defines a model of distributed transaction processing;
- c) defines the requirements to be met by the OSI TP Service; and
- d) takes into consideration the need to coexist with other Application Service Elements, e.g. RDA (Remote Database Access), ROSE (Remote Operations Service Element), and non-ROSE based applications.

This part of ISO/IEC 10026 makes sufficient provisions to allow the specification of transaction-mode communications services and protocols that meet the properties of: atomicity, consistency, isolation, and durability (the ACID properties), as defined in ISO/IEC 9804.

This part of ISO/IEC 10026 does not specify individual implementations or products, nor does it constrain the implementation of entities or interfaces within a computer system.

2 Normative references

The following standards contain provisions which, through reference in this text, constitute provisions of this part of ISO/IEC 10026. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this part of ISO/IEC 10026 are encouraged to investigate the possibility of applying the most recent editions of the standards indicated below. Members of IEC and ISO maintain registers of currently valid International Standards.

ISO/IEC 7498-1:1994, *Information technology - Open Systems Interconnection - Basic Reference Model: The Basic Model.*

ISO 7498-2:1989, *Information processing systems - Open Systems Interconnection - Basic Reference Model - Part 2: Security Architecture.*

ISO/IEC 7498-3:1997, *Information technology - Open Systems Interconnection - Basic Reference Model: Naming and addressing.*

ISO/IEC 8326:1996, *Information technology - Open Systems Interconnection - Session service definition.*

ISO/IEC 8649:1996, *Information technology - Open Systems Interconnection - Service definition for the Association Control Service Element.*

ISO/IEC 8822:1994, *Information technology - Open Systems Interconnection - Presentation service definition.*

ISO/IEC 9545:1989, *Information technology - Open Systems Interconnection - Application Layer structure.*

NOTE - this edition of ISO/IEC 10026 uses the terminology and modelling mechanisms of the first (1989) edition of the Application Layer Structure (ISO/IEC 9545:1989).

ISO/IEC 9579-1:1993, *Information technology - Open Systems Interconnection - Remote Database Access - Part 1: Generic Model, Service, and Protocol.*

ISO/IEC 9594-2:1995, *Information technology - Open Systems Interconnection - The Directory: Models.*

ISO/IEC 9804:1997, *Information technology - Open Systems Interconnection - Service definition for the commitment, concurrency and recovery service element.*

ISO/IEC 10026-2:1998, *Information technology - Open Systems Interconnection - Distributed Transaction Processing - Part 2: OSI TP Service.*

ISO/IEC 10026-3:1998, *Information technology - Open Systems Interconnection - Distributed Transaction Processing - Part 3: Protocol specification.*

ISO/IEC 10026-4:1995, *Information technology - Open Systems Interconnection - Distributed Transaction Processing: Protocol Implementation Conformance Statement (PICS) proforma.*

ISO/IEC 10731:1994, *Information technology - Open Systems Interconnection - Basic Reference Model - Conventions for the definition of OSI services.*

ISO/IEC 13712-1:1995, *Information technology - Remote Operations: Concepts, model and notation.*

3 Definitions

For the purposes of ISO/IEC 10026, the following definitions apply.

3.1 Terms defined in other International Standards

3.1.1 ISO/IEC 10026 makes use of the following terms defined in ISO/IEC 7498-1:

- a) application-entity;
- b) application-process;
- c) application-protocol-data-unit;
- d) concatenation;
- e) open system;
- f) presentation-service;
- g) presentation-service-access-point;
- h) presentation-service-data-unit;
- i) real open system; and
- j) separation.

3.1.2 ISO/IEC 10026 makes use of the following terms defined in ISO 7498-2:

- a) access control;
- b) audit;
- c) authentication;
- d) confidentiality;
- e) integrity; and
- f) non-repudiation.

3.1.3 ISO/IEC 10026 makes use of the following terms defined in ISO/IEC 7498-3:

- a) application-process-invocation-identifier;
- b) application-process-title;
- c) application-entity-invocation-identifier;

- d) application-entity-qualifier; and
- e) application-entity-title.

3.1.4 ISO/IEC 10026 makes use of the following term defined in ISO/IEC 8326:
quality-of-service

3.1.5 ISO/IEC 10026 makes use of the following terms defined in ISO/IEC 10731:

- a) request;
- b) indication;
- c) response;
- d) confirm;
- e) service primitive; primitive;
- f) service-provider; and
- g) service-user.

3.1.6 ISO/IEC 10026 makes use of the following terms defined in ISO/IEC 9545:

- a) application-association; association;
- b) application-context;
- c) application-context-name;
- d) application-entity-invocation;
- e) application-process-invocation;
- f) application-service-element;
- g) association control service element;
- h) multiple association control function;
- i) single association control function; and
- j) single association object.

ITeH STANDARD PREVIEW
(standards.iteh.ai)
ISO/IEC 10026-1:1998
<https://standards.iteh.ai/catalog/standards/sist/8d025d6e-b833-4103-a029-2f1011111111>

3.1.7 ISO/IEC 10026 makes use of the following terms defined in ISO/IEC 9594-2:

- a) Directory Information Tree;
- b) Directory entry; entry;
- c) distinguished name;
- d) object class; and
- e) relative distinguished name.

3.1.8 ISO/IEC 10026 makes use of the following terms defined in ISO/IEC 9804:

- a) atomic action data;
- b) atomicity;
- c) bound data;
- d) consistency;
- e) durability;
- f) final state;
- g) heuristic decision;
- h) initial state; and
- i) isolation.

3.2 Terms defined in ISO/IEC 10026

3.2.1 application-supported distributed transaction: A transaction where the user of the OSI TP Service is responsible for the maintenance of the ACID properties.

3.2.2 chained sequence: A sequence of related contiguous (provider-supported) transaction branches, on the same dialogue, that are aimed at achieving a common goal.

3.2.3 Channel Protocol Machine; CPM: The part of an AEI involved in OSI TP that establishes and terminates TP channels.

3.2.4 channel; Transaction Processing channel: A relationship over an association between two AEIs to facilitate Transaction Processing Service Provider (TPSP) recovery activity. Channels are not visible to the TPSUIs.

3.2.5 commit master: The neighbour to which a node has sent a ready signal.

NOTE - with the static commitment procedures, the commit master will be the dialogue superior.

3.2.6 commit slave: A neighbour from which a ready signal has been received.

NOTE - CCR uses the term "commit subordinate"; TP uses the term "commit slave" to avoid confusion with dialogue subordinate.

NOTE - with the static commitment procedures, a commit slave will be a dialogue subordinate.

NOTE - the terms commit master and commit slave do not apply when a read-only signal or early-exit signal or one-phase signal is sent.

3.2.7 commitment; transaction commitment: Completion of a transaction with the release of transaction data in the final state.

NOTE - commitment requires two-phase commitment procedures if bound data are affected; one-phase commitment procedures may be used if bound data are not affected; see section 8.6.1 for two-phase commitment procedures and 8.6.4 for one-phase commitment procedures.

NOTE - the terms "commitment" and "rollback" have a different scope from that defined in ISO/IEC 9804. ISO/IEC 10026 is concerned with the commitment and rollback of a complete transaction, whereas ISO/IEC 9804 refers to the commitment and rollback of a single atomic action branch.

3.2.8 commitment coordinator: A TPPM involved in a distributed transaction that arbitrates the final outcome of the transaction.

NOTE - With the static two-phase commitment procedures, the commitment coordinator will be at the root of the transaction tree. If the static one-phase commitment procedures are in use in the transaction tree, the commitment coordinator will either be a leaf node or an intermediate node. With the dynamic two-phase commitment procedures, the position of the commitment coordinator may be predetermined or may be determined dynamically.

3.2.9 commitment hinterland: A node's current commitment hinterland is the set of nodes in the transaction tree which include:

- a) the neighbouring nodes from which ready signals have been received; and
- b) the commitment hinterlands of those neighbouring nodes, and so on recursively.

NOTE - the commitment hinterland excludes those nodes which signal read-only or one-phase or early-exit.

NOTE - with the static two-phase commitment procedures and no use of either read-only or one-phase commitment or early-exit, the commitment hinterland of a node will be identical to the transaction subtree of the node.

3.2.10 commitment order: A statement from a node to a neighbour that has signalled ready, that the transaction shall be committed.

3.2.11 control: The permission, on a particular dialogue, for a TPSUI to communicate with its partner.

3.2.12 coordination level: An agreement between two TPSUIs on what mechanism will be used to guarantee the four properties of a transaction; the coordination level may be "commitment", "one-phase commitment" or "none".

3.2.13 coordinated dialogue; dialogue is coordinated: A dialogue currently having a coordination level of "commitment" or "one-phase commitment".

NOTE - a dialogue supporting chained transaction branches is always coordinated and a dialogue supporting unchained transaction branches is coordinated only when it supports a transaction branch.

3.2.14 dialogue: The relationship between two TPSUIs that communicate with each other. The initiator of the dialogue is the superior and the recipient is the subordinate.

3.2.15 dialogue tree: A tree consisting of TPSUIs as the entities with dialogues as the relationships between them.

3.2.16 distributed transaction: A transaction, parts of which may be carried out in more than one open system.

3.2.17 dynamic commitment procedures: The two-phase commit procedures without the constraints of the static commitment procedures; subject to optional controls, the commitment coordinator may be a predetermined node in the transaction tree (not necessarily the root) or may be dynamically determined.

3.2.18 early-exit signal: A statement from a node to a superior that this node and its subtree can make no contribution to the work of the transaction and so it withdraws from participation in the transaction; conditions are that the bound data of this node have not been altered by the transaction, that read-only or early-exit signals have been received from all the node's subordinates in the transaction tree, if there are any, and that reporting of the transaction outcome is not required.

3.2.19 heuristic-hazard: The condition that arises when, as a result of communication failure with a subordinate, the bound data of the subordinate's subtree are in an unknown state.

3.2.20 heuristic-mix: The condition that arises when, as a result of one or more heuristic decisions having been taken, the bound data of the transaction are in an inconsistent state.

3.2.21 intermediate: An entity in a tree which has one superior and one or more subordinates.

3.2.22 leaf: An entity in a tree which has one superior and no subordinates.

3.2.23 local resource: A resource that is resident on the same real open system as the requester of the resource, or a resource that is managed by an entity residing in the same real open system as the requester of the resource.

3.2.24 log-commit record: A record written to the recovery log that reflects the transaction's decision to commit.

3.2.25 log-damage record: A record written to the recovery log that reflects the current inconsistent state of bound data in the subtree.

3.2.26 log-heuristic record: A record written to the recovery log that reflects the node's heuristic decision.

3.2.27 log-ready record: A record written to the recovery log that records information required for recovery and that the bound data of this node is ready-to-commit and, if there is more than one neighbour in the transaction tree, that one of ready signal, one-phase signal or read-only signal or early-exit signal has been received from all but one of the neighbours in the transaction tree.

3.2.28 long lived data: Data which are accessed and manipulated by the TPSUI within the scope of either a provider supported transaction or an application supported transaction but for which the TPSUI takes responsibility for recovery in the event of failures.

NOTE - "long lived data" are not "bound data", and vice versa.

3.2.29 neighbour: An entity in a tree which has a direct relationship with another entity.

NOTE - thus a subordinate and its superior are neighbours, each with the other.

3.2.30 node: A TPSUI together with its TPPM.

3.2.31 node crash: A failure of the node (i.e. TPPM and TPSUI) or of the local environment supporting the node such that dialogues are aborted and all data not recorded in secure storage may be lost.

3.2.32 one-phase signal: A statement from a node to a neighbour that this node has no bound data (in the strict sense defined by CCR) and that either read-only or early-exit or one-phase signals have been received from all other neighbours in the transaction tree, if there are any.

3.2.33 polarized control mode: A mode of communication over a dialogue where only one TPSUI involved in the dialogue is allowed to have control at a time.

3.2.34 Protocol Machine; PM: A generic term to denote either a Transaction Processing Protocol Machine or a Channel Protocol Machine.

3.2.35 provider-supported distributed transaction: A transaction where the provider of the OSI TP Service is responsible for the maintenance of the ACID properties.

3.2.36 read-only signal: A statement from a node to a superior that the bound data of this node have not been altered by the transaction, that read-only or early-exit signals have been received from all the node's subordinates in the transaction tree, if there are any, and that reporting of the transaction outcome is not required.

3.2.37 ready signal: A statement from a node (to a neighbour) that a log-ready record has been written. The neighbour to whom the signal is sent is the one neighbour (if there is more than one) that had not sent a ready signal or one-phase signal or read-only signal or early-exit signal when the log-ready record was written.

NOTE - thus ready signal excludes read-only signal or one-phase signal or early-exit signal.

3.2.38 ready-to-commit state: A state of bound data in which, until the transaction has been terminated by commitment or rollback, the bound data can be released in either their initial or their final state.

3.2.39 recovery: Action taken after a failure to remove undesired consequences of the failure.

3.2.40 recovery log: A repository in secure storage used to record data and state information for the purposes of restart and recovery.

3.2.41 remote resource: A resource that is resident on a different real open system than the real open system making the request for resources.

3.2.42 resource: Data and processing capabilities necessary for a TPSUI to carry out the part of a transaction for which it is responsible.

3.2.43 rollback: transaction rollback: Completion of a transaction with the release of bound data in the initial state.

NOTE - the terms "commitment" and "rollback" have a different scope from that defined in ISO/IEC 9804. ISO/IEC 10026 is concerned with the commitment and rollback of a complete transaction, whereas ISO/IEC 9804 refers to the commitment and rollback of a single atomic action branch.

3.2.44 root: The single entity in a tree which has no superior and has one or more subordinates.

3.2.45 secure storage: A reliable non-volatile place where stored information survives any type of recoverable failure within the real open system.

3.2.46 shared control mode: A mode of communication over a dialogue where both TPSUIs involved in the dialogue have control.

3.2.47 static commitment procedures: The two-phase commitment procedures constrained such that the commit decision is made at the root of the transaction tree and is propagated down the tree.

NOTE - this is equivalent to the commitment procedures of 10026:1992 and 10026:1995.

3.2.48 subordinate: The entity which accepts a relationship (from a superior).

3.2.49 subordinate subtree: The subtree of a subordinate node.

3.2.50 subtree: A subset of a tree. The subtree of a particular node contains

- a) the node itself, called the root node of the subtree; and
- b) the subtrees of each subordinate node of the root node of the subtree, recursively.

A leaf node is its own subtree.

3.2.51 superior: The entity which initiates a relationship.

3.2.52 transaction: A set of related operations characterized by four properties: atomicity, consistency, isolation, and durability. A transaction is uniquely identified by a transaction identifier.

NOTE - For reasons of brevity, the term "transaction" is used as a synonym of the term "provider-supported distributed transaction", from 7.8 onwards.

3.2.53 transaction branch: The portion of a distributed transaction performed by a pair of TPSUIs sharing a dialogue.

NOTE - For reasons of brevity, the term "transaction branch" is used as a synonym of the phrase "branch of provider-supported distributed transaction", from 7.8 onwards.

3.2.54 transaction branch identifier: An unambiguous identifier for a specific branch of a specific transaction.

3.2.55 transaction data: Data which are accessed and manipulated by the TPSUI within the scope of a transaction (either a provider-supported transaction or an application-supported transaction); transaction data is either "bound data" or "long-lived data".

3.2.56 transaction hinterland: The transaction hinterland of node B as viewed from node A is the node B together with the transaction hinterland (as viewed from node B) of all B's neighbouring nodes except A which are participating in or have participated in the current transaction on a transaction branch with B.

NOTE - nodes which are no longer participating in the transaction because they have signalled read-only or early-exit, continue to be part of the transaction hinterland until the transaction is terminated.

3.2.57 transaction identifier: A globally unambiguous identifier for a specific transaction.

3.2.58 transaction logging: The recording of node state information and data in a recovery log.

3.2.59 Transaction Processing Application Service Element; TPASE: That part of a Transaction Processing Protocol Machine (TPPM) which handles the OSI TP Protocol on a single application-association.

3.2.60 Transaction Processing Protocol Machine; TPPM: The provider of the OSI TP Service for exactly one TPSUI. A TPPM handles the OSI TP Protocol on all associations that are used for its TPSUI's activity.

3.2.61 Transaction Processing Service Provider; TPSP: The provider of the OSI TP Service. The TPSP provides the OSI TP Service to all the TPSUIs involved in a particular dialogue tree. The TPSP spans several application-process-invocations (APIs) and is the conceptual view of the OSI TP Service as a whole.

3.2.62 Transaction Processing Service User; TPSU: A user of the OSI TP Service: it refers to a specific set of processing capabilities within an application-process.

3.2.63 TPSU Invocation; TPSUI: A particular instance of a TPSU performing functions for a specific occasion of information processing.

3.2.64 TPSU-title: A name, unambiguous within the scope of the application-process containing the TPSU, which denotes a particular TPSU. The TPSU-title implies the type of processing (capabilities) of this TPSU.

3.2.65 transaction recovery: Action taken after a failure in order to put all the bound data of that transaction into a consistent state.

3.2.66 transaction tree: A tree with nodes as the entities, and transaction branches as the relationship between them.

3.2.67 tree: A set of linked entities arranged in a hierarchical structure and connected by relationships.

3.2.68 unchained sequence: A sequence of non-contiguous (provider-supported) transaction branches, on the same dialogue, that are aimed at achieving a common goal.

3.2.69 uncoordinated dialogue; dialogue is not coordinated: A dialogue currently having a coordination level of "none".

3.2.70 user-ASE: An application-specific ASE.