

---

---

**Information technology — Portable  
Common Tool Environment (PCTE) —**

**Part 1:**  
**Abstract specification**

*Technologies de l'information — Environnement d'outil courant  
portable (PCTE) —*

*Partie 1: Spécifications abstraites*

## Contents

|                                             |           |
|---------------------------------------------|-----------|
| <b>1 Scope</b>                              | <b>1</b>  |
| <b>2 Conformance</b>                        | <b>1</b>  |
| 2.1 Conformance of binding                  | 1         |
| 2.2 Conformance of implementation           | 2         |
| 2.3 Conformance of DDL texts and processors | 3         |
| <b>3 Normative references</b>               | <b>3</b>  |
| <b>4 Definitions</b>                        | <b>4</b>  |
| 4.1 Technical terms                         | 4         |
| 4.2 Other terms                             | 4         |
| <b>5 Formal notations</b>                   | <b>5</b>  |
| <b>6 Overview of PCTE</b>                   | <b>5</b>  |
| 6.1 PCTE structural architecture            | 5         |
| 6.2 Object management system                | 6         |
| 6.3 Object base                             | 6         |
| 6.4 Schema management                       | 6         |
| 6.5 Self-representation and predefined SDSs | 7         |
| 6.6 Object contents                         | 7         |
| 6.7 Process execution                       | 8         |
| 6.8 Monitoring                              | 8         |
| 6.9 Communication between processes         | 8         |
| 6.10 Notification                           | 8         |
| 6.11 Concurrency and integrity control      | 8         |
| 6.12 Distribution                           | 9         |
| 6.13 Replication                            | 9         |
| 6.14 Security                               | 10        |
| 6.15 Accounting                             | 10        |
| 6.16 Implementation limits                  | 10        |
| 6.17 Support of fine-grain objects          | 11        |
| 6.18 Support of object-orientation          | 11        |
| <b>7 Outline of the Standard</b>            | <b>11</b> |

|                                                     |           |
|-----------------------------------------------------|-----------|
| <b>8 Foundation</b>                                 | <b>13</b> |
| 8.1 The state                                       | 13        |
| 8.2 The object base                                 | 14        |
| 8.2.1 Objects                                       | 14        |
| 8.2.2 Attributes                                    | 15        |
| 8.2.3 Links                                         | 16        |
| 8.3 Types                                           | 17        |
| 8.3.1 Object types                                  | 17        |
| 8.3.2 Attribute types                               | 18        |
| 8.3.3 Link types                                    | 19        |
| 8.3.4 Enumeral types                                | 23        |
| 8.4 Types in SDS                                    | 23        |
| 8.4.1 Object types in SDS                           | 25        |
| 8.4.2 Attribute types in SDS                        | 25        |
| 8.4.3 Link types in SDS                             | 26        |
| 8.4.4 Enumeral types in SDS                         | 26        |
| 8.5 Types in working schema                         | 26        |
| 8.5.1 Object types in working schema                | 27        |
| 8.5.2 Attribute types in working schema             | 28        |
| 8.5.3 Link types in working schema                  | 28        |
| 8.5.4 Enumeral types in working schema              | 28        |
| 8.6 Types in global schema                          | 28        |
| 8.7 Operations                                      | 29        |
| 8.7.1 Calling process                               | 29        |
| 8.7.2 Direct and indirect effects                   | 29        |
| 8.7.3 Errors                                        | 32        |
| 8.7.4 Operation serializability                     | 33        |
| <b>9 Object management</b>                          | <b>33</b> |
| 9.1 Object management concepts                      | 33        |
| 9.1.1 The basic type "object"                       | 33        |
| 9.1.2 The common root                               | 37        |
| 9.1.3 Datatypes for object management               | 37        |
| 9.2 Link operations                                 | 37        |
| 9.3 Object operations                               | 47        |
| 9.4 Version operations                              | 61        |
| <b>10 Schema management</b>                         | <b>69</b> |
| 10.1 Schema management concepts                     | 69        |
| 10.1.1 Schema definition sets and the SDS directory | 69        |
| 10.1.2 Types                                        | 70        |
| 10.1.3 Object types                                 | 71        |
| 10.1.4 Attribute types                              | 72        |

|               |                                               |            |
|---------------|-----------------------------------------------|------------|
| <b>10.1.5</b> | <b>Link types</b>                             | <b>73</b>  |
| <b>10.1.6</b> | <b>Enumeral types</b>                         | <b>74</b>  |
| <b>10.1.7</b> | <b>Datatypes for schema management</b>        | <b>75</b>  |
| <b>10.2</b>   | <b>SDS update operations</b>                  | <b>75</b>  |
| <b>10.3</b>   | <b>SDS usage operations</b>                   | <b>105</b> |
| <b>10.4</b>   | <b>Working schema operations</b>              | <b>112</b> |
| <b>11</b>     | <b>Volumes, devices, and archives</b>         | <b>117</b> |
| <b>11.1</b>   | <b>Volume, device, and archiving concepts</b> | <b>117</b> |
| <b>11.1.1</b> | <b>Volumes</b>                                | <b>117</b> |
| <b>11.1.2</b> | <b>Administration volumes</b>                 | <b>118</b> |
| <b>11.1.3</b> | <b>Devices</b>                                | <b>118</b> |
| <b>11.1.4</b> | <b>Archives</b>                               | <b>119</b> |
| <b>11.2</b>   | <b>Volume, device, and archive operations</b> | <b>120</b> |
| <b>12</b>     | <b>Files, pipes, and devices</b>              | <b>128</b> |
| <b>12.1</b>   | <b>File, pipe, and device concepts</b>        | <b>128</b> |
| <b>12.2</b>   | <b>File, pipe, and device operations</b>      | <b>132</b> |
| <b>13</b>     | <b>Process execution</b>                      | <b>140</b> |
| <b>13.1</b>   | <b>Process execution concepts</b>             | <b>140</b> |
| <b>13.1.1</b> | <b>Static contexts</b>                        | <b>140</b> |
| <b>13.1.2</b> | <b>Foreign execution images</b>               | <b>141</b> |
| <b>13.1.3</b> | <b>Execution classes</b>                      | <b>141</b> |
| <b>13.1.4</b> | <b>Processes</b>                              | <b>142</b> |
| <b>13.1.5</b> | <b>Initial processes</b>                      | <b>149</b> |
| <b>13.1.6</b> | <b>Profiling and monitoring concepts</b>      | <b>150</b> |
| <b>13.2</b>   | <b>Process execution operations</b>           | <b>150</b> |
| <b>13.3</b>   | <b>Security operations</b>                    | <b>166</b> |
| <b>13.4</b>   | <b>Profiling operations</b>                   | <b>171</b> |
| <b>13.5</b>   | <b>Monitoring operations</b>                  | <b>172</b> |
| <b>14</b>     | <b>Message queues</b>                         | <b>174</b> |
| <b>14.1</b>   | <b>Message queue concepts</b>                 | <b>174</b> |
| <b>14.2</b>   | <b>Message queue operations</b>               | <b>177</b> |
| <b>15</b>     | <b>Notification</b>                           | <b>183</b> |
| <b>15.1</b>   | <b>Notification concepts</b>                  | <b>183</b> |
| <b>15.1.1</b> | <b>Access events and notifiers</b>            | <b>183</b> |
| <b>15.1.2</b> | <b>Notification messages</b>                  | <b>184</b> |
| <b>15.1.3</b> | <b>Time of sending notification messages</b>  | <b>185</b> |
| <b>15.1.4</b> | <b>Range of concerned message queues</b>      | <b>185</b> |
| <b>15.2</b>   | <b>Notification operations</b>                | <b>185</b> |

|                                                            |            |
|------------------------------------------------------------|------------|
| <b>16 Concurrency and integrity control</b>                | <b>187</b> |
| 16.1 Concurrency and integrity control concepts            | 187        |
| 16.1.1 Activities                                          | 187        |
| 16.1.2 Resources and locks                                 | 190        |
| 16.1.3 Lock modes                                          | 192        |
| 16.1.4 Inheritance of locks                                | 194        |
| 16.1.5 Establishment and promotion of locks                | 195        |
| 16.1.6 Implied locks                                       | 196        |
| 16.1.7 Conditions for establishment or promotion of a lock | 197        |
| 16.1.8 Releasing locks                                     | 198        |
| 16.1.9 Permanence of updates                               | 199        |
| 16.1.10 Tables for locks                                   | 200        |
| 16.2 Concurrency and integrity control operations          | 202        |
| <b>17 Replication</b>                                      | <b>208</b> |
| 17.1 Replication concepts                                  | 208        |
| 17.1.1 Replica sets                                        | 208        |
| 17.1.2 Replicated objects                                  | 209        |
| 17.1.3 Selection of an appropriate replica                 | 210        |
| 17.1.4 Administration replica set                          | 211        |
| 17.2 Replication operations                                | 212        |
| <b>18 Network connection</b>                               | <b>218</b> |
| 18.1 Network connection concepts                           | 218        |
| 18.1.1 Execution sites                                     | 218        |
| 18.1.2 Workstations                                        | 219        |
| 18.1.3 Foreign systems                                     | 222        |
| 18.1.4 Network partitions                                  | 222        |
| 18.1.5 Accessibility                                       | 223        |
| 18.1.6 Workstation closedown                               | 225        |
| 18.2 Network connection operations                         | 226        |
| 18.3 Foreign system operations                             | 231        |
| 18.4 Time operations                                       | 233        |
| <b>19 Discretionary security</b>                           | <b>234</b> |
| 19.1 Discretionary security concepts                       | 234        |
| 19.1.1 Security groups                                     | 234        |
| 19.1.2 Access control lists                                | 238        |
| 19.1.3 Discretionary access modes                          | 241        |
| 19.1.4 Access control lists on object creation             | 243        |
| 19.2 Operations for discretionary access control operation | 244        |
| 19.3 Discretionary security administration operations      | 248        |

|                                                   |            |
|---------------------------------------------------|------------|
| <b>20 Mandatory security</b>                      | <b>253</b> |
| 20.1 Mandatory security concepts                  | 253        |
| 20.1.1 Mandatory classes                          | 253        |
| 20.1.2 The mandatory class structure              | 255        |
| 20.1.3 Labels and the concept of dominance        | 256        |
| 20.1.4 Mandatory rules for information flow       | 258        |
| 20.1.5 Multi-level security labels                | 261        |
| 20.1.6 Floating security levels                   | 264        |
| 20.1.7 Implementation restrictions                | 266        |
| 20.1.8 Built-in policy aspects                    | 266        |
| 20.2 Operations for mandatory security operation  | 268        |
| 20.3 Mandatory security administration operations | 274        |
| 20.4 Mandatory security operations for processes  | 279        |
| <b>21 Auditing</b>                                | <b>281</b> |
| 21.1 Auditing concepts                            | 281        |
| 21.1.1 Audit files                                | 281        |
| 21.1.2 Audit selection criteria                   | 283        |
| 21.2 Auditing operations                          | 284        |
| <b>22 Accounting</b>                              | <b>289</b> |
| 22.1 Accounting concepts                          | 289        |
| 22.1.1 Consumers and accountable resources        | 289        |
| 22.1.2 Accounting logs and accounting records     | 290        |
| 22.2 Accounting administration operations         | 294        |
| 22.3 Consumer identity operations                 | 299        |
| <b>23 Common binding features</b>                 | <b>300</b> |
| 23.1 Mapping of types                             | 300        |
| 23.1.1 Mapping of predefined PCTE datatypes       | 300        |
| 23.1.2 Mapping of designators and nominators      | 302        |
| 23.1.3 Mapping of other values                    | 310        |
| 23.2 Object reference operations                  | 311        |
| 23.3 Link reference operations                    | 314        |
| 23.4 Type reference operations                    | 317        |
| <b>24 Implementation limits</b>                   | <b>320</b> |
| 24.1 Bounds on installation-wide limits           | 320        |
| 24.2 Bounds on workstation-dependent limits       | 321        |
| 24.3 Limit operations                             | 322        |
| 24.3.1 Datatypes for limit operations             | 322        |

|                                                                            |            |
|----------------------------------------------------------------------------|------------|
| <b>Annex A - VDM Specification Language for the Abstract Specification</b> | <b>323</b> |
| <b>Annex B - The Data Definition Language (DDL)</b>                        | <b>329</b> |
| <b>Annex C - Specification of Errors</b>                                   | <b>339</b> |
| <b>Annex D - Auditable Events</b>                                          | <b>362</b> |
| <b>Annex E - The Predefined Schema Definition Sets</b>                     | <b>370</b> |
| <b>Annex F - The fine-grain objects module</b>                             | <b>387</b> |
| <b>Annex G - The object-orientation module</b>                             | <b>400</b> |
| <b>Index of Operations</b>                                                 | <b>425</b> |
| <b>Index of Error Conditions</b>                                           | <b>431</b> |
| <b>Index of Technical Terms</b>                                            | <b>439</b> |

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 13719-1:1998

## 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 13719-1 was prepared by ECMA (as Standard ECMA-149) and was adopted, under a special “fast-track procedure”, by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, in parallel with its approval by national bodies of ISO and IEC.

This second edition cancels and replaces the first edition (ISO/IEC 13719-1:1995), which has been technically revised.

ISO/IEC 13719 consists of the following parts, under the general title *Information technology - Portable Common Tool Environment (PCTE)*:

- Part 1: Abstract specification
- Part 2: C programming language binding
- Part 3: Ada programming language binding
- Part 4: IDL binding (Interface Definition Language)

Annexes A to D and annexes F and G form an integral part of this part of ISO/IEC 13719. Annex E is for information only.



# Information technology — Portable Common Tool Environment (PCTE) —

## Part 1:

### Abstract specification

## 1 Scope

This part of ISO/IEC 13719 specifies PCTE in abstract, programming-language-independent, terms. It specifies the interface supported by any conforming implementation as a set of abstract operation specifications, together with the types of their parameters and results. It is supported by a number of standard *bindings*, i.e. representations of the interface in standard programming languages.

The scope of this part of ISO/IEC 13719 is restricted to a single PCTE installation. It does not specify the means of communication between PCTE installations, nor between a PCTE installation and another system.

A number of features are not completely defined in this part of ISO/IEC 13719, some freedom being allowed to the implementor. Some of these are *implementation limits*, for which constraints are defined (see clause 24). The other implementation-dependent and implementation-defined features are specified in the appropriate places in this Standard.

PCTE is an interface to a set of facilities that forms the basis for constructing environments supporting systems engineering projects. These facilities are designed particularly to provide an infrastructure for programs which may be part of such environments. Such programs, which are used as aids to systems development, are often referred to as tools.

This part of ISO/IEC 13719 also includes (in annex B) a language standard for the PCTE Data Description Language (DDL), suitable for writing PCTE schema definition sets.

## 2 Conformance

### 2.1 Conformance of binding

A binding conforms to this part of ISO/IEC 13719 if and only if:

- it consists of a set of operational interfaces and datatypes, with a mapping from the operations and datatypes of this part of ISO/IEC 13719;
- each operation of this part of ISO/IEC 13719 is mapped to one or more sequences of one or more operations of the binding (distinct operations need not be mapped to distinct sets of sequences of binding operations);
- each datatype of this part of ISO/IEC 13719 is mapped to one or more datatypes of the binding;
- each named error of this part of ISO/IEC 13719 is mapped to one or more error values (status values, exceptions, or the like) of the binding;
- the conditions of clause 23 on common binding features are satisfied;

- the conditions for conformance of an implementation to the binding are defined, are achievable, and are not in conflict with the conditions in 2.2 below.

## 2.2 Conformance of implementation

The functionality of PCTE is divided into the following modules:

- The core module consists of the datatypes and operations defined in clauses 8 to 19 (except 13.1.6, 13.4, and 13.5) and 23.
- The mandatory access control module consists of the datatypes and operations defined in clause 20.
- The auditing module consists of the datatypes and operations defined in clause 21.
- The accounting module consists of the datatypes and operations defined in clause 22.
- The profiling module consists of the datatypes defined in 13.1.6 and the operations defined in 13.4.
- The monitoring module consists of the datatype Address defined in 13.1.6 and operations defined in 13.5.
- The fine-grain objects module consists of the following extensions defined in annex F:
  - . extensions to the semantics of operations to cater for fine-grain objects;
  - . new operations;
  - . new error conditions;
  - . additions to the predefined SDS system.
- The object-orientation module consists of the following extensions defined in annex G:
  - . additions to the predefined SDSs metasds and system;
  - . an extension to the semantics of the operation SDS\_REMOVE\_TYPE to cater for the new classes of type;
  - . new operations;
  - . new error conditions.

An implementation of PCTE conforms to this part of ISO/IEC 13719 if and only if it implements the core module.

An implementation of PCTE conforms to this part of ISO/IEC 13719 with mandatory access control level 1 or 2 if it implements the core module and in addition:

- for level 1: the mandatory access control module except the floating security levels features defined in 20.1.6;
- for level 2: the mandatory access control module.

An implementation of PCTE conforms to this part of ISO/IEC 13719 with auditing if and only if it implements the core module and in addition the auditing module.

An implementation of PCTE conforms to this part of ISO/IEC 13719 with accounting if and only if it implements the core module and in addition the accounting module.

An implementation of PCTE conforms to this part of ISO/IEC 13719 with profiling if and only if it implements the core module and in addition the profiling module.

An implementation of PCTE conforms to this part of ISO/IEC 13719 with monitoring if and only if it implements the core module and in addition the monitoring module.

An implementation of PCTE conforms to this part of ISO/IEC 13719 with fine-grain objects if and only if it implements the core module and in addition, implements the fine-grain objects module.

An implementation of PCTE conforms to this part of ISO/IEC 13719 with object-orientation if and only if it implements the core module and in addition the object-orientation module.

By 'an implementation implements a module' is meant that, for the clauses of the module:

- the implementation conforms to a binding of this part of ISO/IEC 13719 which itself conforms to this part of ISO/IEC 13719 and which is itself an International Standard;
- if an operation of this part of ISO/IEC 13719 is mapped to a set of sequences of operations in the binding:
  - . case 1: operation\_A; operation\_B; ... operation\_F;
  - . case 2: operation\_G; operation\_H; ...operation\_M;
  - . etc.

then in each case the sequence of invocations of the operations of the implementation must have the effect of the original operation of this part of ISO/IEC 13719;

- the relevant limits on quantities specified in clause 24 are no more restrictive than the values specified there;
- the implementations of the implementation-defined features in this part of ISO/IEC 13719 are all defined.

An implementation of PCTE does not conform to this part of ISO/IEC 13719 if it implements any of the following, whether or not the PCTE entity mentioned is in a module which the implementation implements:

- an operation with same name as a PCTE operation but with different effect;
- an SDS with the same name as a PCTE predefined SDS but with different contents;
- an error condition with the same name as a PCTE error condition but with different meaning.

### 2.3 Conformance of DDL texts and processors

A DDL definition conforms to this part of ISO/IEC 13719 if it conforms to the syntax and obeys the constraints of the DDL definition in annex B.

A DDL processor conforms to this part of ISO/IEC 13719 if it accepts any conforming DDL definition and processes it in conformance with the meaning of DDL as defined in annex B.

## 3 Normative references

The following standards contain provisions which, through reference in this text, constitute provisions of this part of ISO/IEC 13719. 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 13719 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 2022:1994,    | <i>Information technology - Character code structure and extension techniques.</i>                                                                                                      |
| ISO 8601:1988,        | <i>Data elements and interchange formats - Information interchange - Representation of dates and times.</i>                                                                             |
| ISO/IEC 8859-1:1998,  | <i>Information technology - 8-bit single-byte coded graphic character sets - Part 1: Latin alphabet No. 1.</i>                                                                          |
| ISO/IEC 10646-1:1993, | <i>Information technology - Universal Multiple-Octet Coded Character Set (UCS) - Part 1: Architecture and Basic Multilingual Plane.</i>                                                 |
| ISO/IEC 11404:1996,   | <i>Information technology - Programming languages, their environments and system software interfaces - Language-independent datatypes.</i>                                              |
| ISO/IEC 13817-1:1996, | <i>Information technology - Programming languages, their environments and system software interfaces - Vienna Development Method - Specification Language - Part 1: Basic language.</i> |
| ISO/IEC 14977:1996,   | <i>Information technology - Syntactic metalanguage - Extended BNF.</i>                                                                                                                  |

## 4 Definitions

### 4.1 Technical terms

All technical terms used in this part of ISO/IEC 13719, other than a few in widespread use, are defined in the text, usually in a formal notation. All identifiers defined in VDM-SL or in DDL (see clause 5) are technical terms; apart from those, a defined technical term is printed in italics at the point of its definition, and only there. For the use of technical terms defined in VDM-SL and DDL see clause A.3 and clause B.9 respectively. All defined technical terms are listed in an index, with references to their definitions.

### 4.2 Other terms

For the purposes of this part of ISO/IEC 13719, the following definitions apply.

**4.2.1 implementation-defined:** Possibly differing between PCTE implementations, but defined for any particular PCTE implementation.

**4.2.2 implementation-dependent:** Possibly differing between PCTE implementations and not necessarily defined for any particular PCTE implementation.

**4.2.3 binding-defined:** Possibly differing between language bindings, but defined for any particular language binding.

**4.2.4 datatype:** The type of a parameter or result of an operation defined in this part of ISO/IEC 13719, or used to define such a type. Where, as in clause 23, it is necessary to distinguish these types from datatypes defined elsewhere, the term *PCTE datatype* is used.

**4.2.5 operation:** a name plus a signature that is used in the context of an invocation to trigger the execution of a specific method.

**4.2.6 interface:** a set of operations; interfaces are a convenient way to group operations so that they can be referred to together, e.g. to define other interfaces by inheritance.

**4.2.7 method:** the set of actions triggered by an operation.

## 5 Formal notations

Four formal notations are used in this part of ISO/IEC 13719.

For datatypes and for operation signatures, a small subset of the *Vienna Development Method Specification Language* or *VDM-SL* is used; it is defined in annex A. This subset of VDM-SL is also used to define some types used for operation parameters and results.

The *Data Definition Language* or *DDL* is used to define types; it is defined in annex B. Where a concept is defined in both VDM-SL and DDL, the same identifier is used.

To define the error conditions detected by operations, a parameterized notation is used; it is defined in annex C.

The BSI syntactic notation (BS 6154 : 1981) is used to define the syntax of VDM-SL and DDL, and in a few other places where the syntax of strings is defined.

## 6 Overview of PCTE

PCTE is designed to support program portability by providing machine-independent access to a set of facilities. These facilities, which are described in this part of ISO/IEC 13719, are designed particularly to provide an infrastructure for programs to support systems engineering projects.

The PCTE architecture is described in two dimensions: the *structural architecture* and the *functional architecture*. The structural architecture is described in 6.1, and shows how a PCTE installation is built of a system of communicating workstations and how the software providing the PCTE interfaces is structured. The functional architecture is described in 6.2 onwards, and gives an outline of the functional components of PCTE and the facilities they provide.

### 6.1 PCTE structural architecture

The preferred structural architecture for a PCTE installation is a set of workstations and associated resources communicating over a network, though other architectures are possible. There is no hierarchy or ordering of workstations within a PCTE installation. If a workstation is part of a PCTE installation then the PCTE installation appears to the workstation's user as a conceptually single machine, although each workstation can act as an autonomous unit. Such a user has access to the total resources of a PCTE installation, subject to the necessary access controls.

The PCTE database (called the *object base*) is partitioned into volumes. Volumes are dynamically allocated to (*mounted on*) particular workstations, and, once mounted, are globally available in that PCTE installation.

The program writer does not need to be aware of the distribution architecture, but the PCTE interfaces do provide all the facilities needed to configure a PCTE installation and control its distribution. The PCTE interfaces appear to the tool writer as available within a PCTE installation irrespective of the tool's physical location within a PCTE installation and independent of any particular network topology.

## 6.2 Object management system

An aspect of PCTE that is of major importance to the process of constructing and integrating portable tools is the provision of the object base and a set of functions to manipulate the various objects in the object base. The object base is the repository of the data used by the tools of a PCTE installation, and the *Object Management System* or *OMS* of PCTE provides the functions used to access the object base.

In a general sense, the users and programs of the PCTE installation have the ability to manage entities that are known to, and can be designated in, a particular PCTE installation. These may be files in the traditional sense, or peripherals, interprocess message queues or pipes, or the description of processes themselves or of the static context of a process. Tools supporting user applications establish classes of objects defined by the user: these can represent information items such as project milestones, tasks, and change requests.

## 6.3 Object base

The basic OMS model is derived from the Entity Relationship data model and defines *objects* and *links* as being the basic items of a PCTE object base.

Objects are entities (in the Entity Relationship sense) which can be designated, and can optionally have:

- *Contents*: a storage of data representing the traditional file concept;
- *Attributes*: primitive values representing specific properties of an object which can be named individually;
- *Links*: representations of associations between objects. Links may have attributes, which may be used to describe properties of the associations or as keys to distinguish between links of the same type from the same object.

Designation of links is the basis for the designation of objects: the principal means for accessing objects in most OMS operations is to navigate the object base by traversing a sequence of links.

## 6.4 Schema management

Entities used by the user and those used by the system that are represented by objects in the object base can be treated in a uniform manner, and facilities to control their structure, to store and to designate these objects, are provided by PCTE.

The object base of each PCTE installation is governed by a typing mechanism. All entities in the object base are typed and the data must conform to the corresponding type rules. Type rules are defined for objects, for links, and for attributes.

PCTE is designed to allow, but not to require, distributed and devolved management of the object base. To this end the definition of the typing rules which govern an object, a link, or an attribute in the object base may be split up among a number of *schema definition sets* (or *SDSs*). Some properties of an object, a link, or an attribute must be the same in every SDS which contributes to the definition of the typing rules for that object, link, or attribute: these are properties of the *type*. Other properties may differ for different SDSs: these are properties of the *type in SDS*.

Each SDS provides a consistent and self-contained view of the data in the object base. A process, at any one time, views the data in the object base through a *working schema*. A working schema is obtained as a composition of SDSs in an ordered list. The effect of such a composition is to provide a union of all the types contained in the listed SDSs. A uniform naming algorithm, dependent on the ordering of the SDSs, is applied to all the contained types.

The object base of a PCTE installation has a notional *global schema*, composed of all the SDSs. The global schema is not directly represented in the object base, and the concept is used mainly to state certain consistency constraints on the object base as a whole.

Child types of object types can be defined with the effect of implicit inheritance of all properties of their parent types. Additionally, child types can have properties of their own.

## 6.5 Self-representation and predefined SDSs

Many of the entities in a PCTE installation are represented by objects in the object base. The types of these objects are defined in *predefined SDSs*, which are available in any conforming implementation; for example processes are represented by objects of type "process" which is defined in the predefined SDS 'system'. This property of PCTE is called *self-representation*. In general, in this part of ISO/IEC 13719, the name of an entity is used also to refer to the object that represents it.

In some cases an object of a type representing some kind of entity requires initializing, or must be created by a particular operation, before it can be used in operations to represent an entity of that kind. Such an object which has been initialized or correctly created is referred to as a *known* entity of that kind (i.e. known to the PCTE installation); any other object of that type is referred to as an *unknown* entity. For example an object of type "process" created by PROCESS\_CREATE is a known process, while one created by OBJECT\_CREATE is an unknown process.

## 6.6 Object contents

A set of operations is provided to access the contents of some types of objects (files, pipes, and devices). These operations provide conventional input-output facilities on files and pipes and control of input and output on devices. These contents are not interpreted by PCTE.

Other types of objects (accounting logs and audit files) have contents with structure that is defined by PCTE and for access to which special operations are provided.



## 6.7 Process execution

PCTE is an interface to support programs. When a program is *run*, this is either the *execution* of the program itself, or the execution of an interpreter which interprets the program. An execution of a program is a *process*. Processes are represented by objects in the object base, so the hierarchy of processes, the environment in which a process runs, the parameters it has been passed, and the various stages of the program execution can be controlled, manipulated and examined.

These facilities can be used also to control processes running on *foreign systems*. A *foreign system* can be a foreign development system, a target system running a real-time operating system, or even a PCTE workstation in another PCTE installation.

## 6.8 Monitoring

PCTE provides three sets of features to support debugging and monitoring of processes.

- To measure the amount of time spent in selected parts of the code.
- To observe, and modify, the execution of a child process.
- To measure the processor usage of the calling process.

## 6.9 Communication between processes

PCTE provides a number of different mechanisms for communicating between processes. The principal ones supplied are:

- the objects, links and attributes in the database;
- message queues;
- pipes.

Message queues and pipes are essentially special forms of object. Thus both pipes and message queues are special cases of the general use of the object base for *interprocess communication*.

Pipes and message queues also provide communication between PCTE processes and foreign processes running on foreign systems (if the foreign systems allow it).

## 6.10 Notification

In PCTE there is a mechanism that allows the designation of objects so that certain types of access result in a message being posted in a message queue which can be accessed by the process requesting the notification.

The notification mechanism allows a process to specify events, corresponding to operations on objects, of which it wants to be notified.

## 6.11 Concurrency and integrity control

The object base is subject to concurrent access by users, and is liable to underlying system failure.



PCTE provides locking facilities to control the strength of object base concurrency and consistency, ranging from unprotected behaviour, through protected behaviour, to protected atomic and serializable transaction activities. PCTE ensures object base consistency and object base integrity for atomic and serializable transactions.

Each user carrying out a transaction on the object base sees some grouping of operations as an atomic operation which transforms the object base from one consistent state to another. If transactions are run one at a time then each transaction sees the consistent state left by its predecessor. When transactions are run concurrently PCTE ensures that the effect on the object base is as though they were run serially. With a few exceptions, such as messages sent to or received from a message queue, the effect of a sequence of operations performed within a transaction is atomic: either all the operations are performed or none are performed.

Another important aspect of activities arises in composition of programs. A single program carrying out an atomic transaction on the object base can be regarded as performing a single function. More powerful functions can be built up by an outer program invoking a set of other, inner, programs, each of which carries out its own specific function. PCTE provides *nested activities* to allow each inner activity to behave in an atomic way, and at the same time to allow the whole function to be atomic. Thus the outer program can start a transaction, which may be either committed or aborted, and finally the whole outer transaction is committed or aborted. Each such inner program could itself invoke further nested programs, and so on.

## 6.12 Distribution

PCTE is based on a community of workstations of possibly differing types connected together by a network. The community is normally seen by the user as a single environment, grouping together the facilities, services and resources of all the different workstations, though in some circumstances a PCTE installation may be temporarily divided into separated partitions, each of which supports useful work.

Objects, including processes, are distributed throughout a PCTE installation. A user is able to disregard both the location of objects on volumes in the network and that of the workstation concerned in executing processes. Alternatively a user may choose to exercise control over the location of objects on volumes and the location of processes. On creation of an object a volume can be specified to indicate its location. Every process executes on a particular workstation and a user can specify which workstation by either static or dynamic means: the static context of a program has an execution class identifying the range of workstations upon which the static context may be executed; the workstation on which a process executes can be specified on invocation.

## 6.13 Replication

As it is possible that one or more workstations of a PCTE installation become temporarily unavailable, certain installation-wide objects must still be accessible. Replication facilities are available whereby a copy of an object's contents, attributes and links are made to each workstation. Installation-wide objects are predefined as replicated and other objects can be added. This feature is intended for non-volatile, rarely varying, widely consulted objects.

## 6.14 Security

A PCTE installation has to support many users and many projects. Different users are expected to have different roles within projects and to be authorized to access different objects. The user accesses objects using programs (themselves modelled as static contexts within the object base).

The purpose of security is to prevent the unauthorized disclosure, amendment or deletion of information. Security facilities are provided to support the definition of the different authorizations of users and programs.

Security in PCTE is provided by discretionary and mandatory access controls. Access controls as defined in the security clauses form one aspect of the correct operation of the installation with regard to the integrity of the information held and the correctness of its use. In this regard, the facilities described in the security clauses complement the data modelling facilities of the OMS and schema management, and the transaction and concurrency control facilities.

Each OMS object is associated with *access control lists* which define which types of access to the object are permitted for designated users or programs. Access control lists are expressed in terms of *discretionary access rights* which are explicitly granted or denied to designated individual users, user groups or program groups. Access rights on a particular object are combined in order to determine a process's permission to perform each particular operation on the object.

Mandatory access controls cover both *mandatory confidentiality* and *mandatory integrity*, with distinct controls. Mandatory access controls are additional to discretionary access controls.

Mandatory confidentiality controls prevent the disclosure of information to unauthorized users. They prevent the flow of information to the unauthorized user directly, by controlling read access (*simple confidentiality*), and indirectly, by controlling the flow of information between objects (*confidentiality confinement*).

Mandatory integrity controls prevent unauthorized sources from contributing to the information in an object. They prevent the flow of information from the unauthorized user directly, by controlling write access (*simple integrity*), and indirectly, by controlling the flow of information between objects (*integrity confinement*).

## 6.15 Accounting

The accounting facilities of PCTE allow the automatic recording of the consumption of selected installation resources by users, groups of users, or groups of programs.

Authorized users may designate selected objects like programs, files, pipes, message queues, devices, workstations, and SDSs as being accountable resources. Access to an accountable resource by a process implies the automatic logging of usage information into the associated accounting log on completion of the operation.

## 6.16 Implementation limits

PCTE permits the user to examine the implementation-defined limits for the PCTE installation in which a program executes.

Minimal values are defined for limits, so that a program respecting those values is portable to any PCTE installation.

### 6.17 Support of fine-grain objects

The notion of support of fine-grain objects is mostly concerned with improved performance time for creating and accessing PCTE objects. Object granularity is not dependent on type. It is described in terms of the amount of processing that has to be done to access an object.

To enhance performance, the concept of cluster is introduced. A *cluster* is an object that represents the set of fine-grain objects that share the same values for certain PCTE properties and with some specific restrictions:

- Usage restrictions on concurrency allow them to be cached in the main memory of processes.
- Time attributes of all fine-grain objects residing in a cluster are shared.
- Notification is not applicable to fine-grain objects.
- Security properties are also shared and only checked once at the level of the cluster.
- Auditing has limitations which decreases the controls to be made on fine-grain objects.
- Fine-grain objects are not accountable resources.
- Fine-grain objects have the same replicated state as their cluster.

### 6.18 Support of object-orientation

One of the prominent characteristics of PCTE is its ability to define any user data model and to use a self-referential approach to describe its metadata. The object-orientation facilities follow a similar approach and describe everything as an extension of the metabase or of the object base.

The data model supporting the object-orientation facilities can be partitioned into three parts: the interface part, the module part, and the method mapping part.

The interface part of this data model is described as an extension of the metasds SDS, while the other two parts are extensions of the system SDS. The reason is that the instances of the interface part are additional information contained in a user SDS, while the instances of the two other parts are user data stored in the object base, as executable programs or loadable modules.

## 7 Outline of the Standard

Clause 6 gives an informal, non-normative explanation of the concepts of PCTE. Clause 7 gives an overview of the document and of the structure of the definition.

The partly formal, normative definition of PCTE is in clauses 8 to 24 and annexes A to C. It is in two main parts. The first main part is the *foundation* (clause 8) which defines the concept Object and its parts, for example Attribute and Link, and the concepts of the associated typing mechanism, for example Type and Type in SDS. This uses a subset of VDM-SL; see annex A.

The second main part of the definition is the interface definition (clauses 9-22). This defines the other concepts of PCTE, for example Process and Workstation, as specializations of the concept Object (clauses 11-22). This definition is in terms of the typing structure associated with these specializations, that is in terms of the typing concepts of the foundation. A language for the definition of types and types in SDS, called *Data Definition Language* or *DDL*, is defined in annex B.

The concept Object is itself further specialized, i.e. details not necessary for the foundation are added, in clause 9. (The name *Object* is used in both the foundation and the interface definition because it is the same concept although only a few of its details are defined in the foundation.)

Thus the foundation is a relatively simple general model that is specialized in later clauses to provide the PCTE interface definition.

Instances of the PCTE concepts are called *entities* and they are referred to by the names of the underlying concepts, for example instances of Object are called objects. All the entities existing at a time are called the *state* of the PCTE installation. PCTE is defined in terms of the permissible values of the state and the permissible operations on the state. The foundation defines part of the state, namely that part concerned with entities of the foundation concepts; the interface definition defines the rest of the state and all the operations.

The concepts of the typing mechanism cannot be treated as specializations of the concept Object because the definition of PCTE would then be circular. They can however be *represented* by specializations of Object so that tools can determine the current state of the typing mechanism using the operations provided for determining the current state of objects. Operations for manipulating the state of the typing mechanism also manipulate the representing objects automatically and equivalently. The representations and operations of the typing mechanism are defined in clause 10.

The interface is defined by operations grouped according to function. For each group some concepts are defined first in DDL and possibly VDM-SL, as described above. There follow the operation definitions; a VDM-SL definition of the signature, an informal English description of the normal action of the operation, and a list of the possible error conditions (using an abbreviated notation defined in annex C).

Other parts of this International Standard define application programming interfaces to PCTE in terms of specific programming languages by defining the mapping of datatypes, operations, and error conditions of the abstract specification to datatypes, operations, and error conditions respectively of the programming language (see 3.1). Such mapping specifications are called *bindings*. Clause 23 defines a number of features to which all bindings must conform.

Clause 24 defines the limits on the sizes and numbers of various entities which a conforming PCTE implementation must respect. These are given as minima which an implementation must meet or exceed.

Annexes A to C define various notations used in the Abstract Specification. Annex A defines the subset of VDM-SL used for type definitions and operation signatures; annex B defines DDL; and annex C defines the notation for operation error conditions.

Annex D contains a list of auditable events classified by event type.

Annex E is provided for information; it collects the DDL definitions of the types in the predefined schema definition sets.

Annex F is normative and contains the definition of the fine-grain objects module.

Annex G is normative and contains the definition of the object-orientation module.

Clauses 8 to 24 contain commentary (headed NOTE or NOTES) which is not normative and is intended as a help to the reader in understanding the definition.

## 8 Foundation

### 8.1 The state

```

state PCTE_Installation of
  SYSTEM_TIME      : Time
  OBJECT_BASE      : map Object_designator to Object
  PROCESSES        : set of Process
  MESSAGE_QUEUES   : set of Message_queue
  CONTENTS_HANDLES : map Contents_handle to Current_position
  CURRENT_POSITIONS : map Current_position to Natural
  WORKSTATIONS     : set of Workstation
end

Name = Text

Name_sequence = seq of Name

Working_schema ::
  VISIBLE_TYPES : set of Type_in_working_schema
  SDS_NAMES     : Name_sequence

Process ::
  PROCESS_OBJECT      : Object_designator
  WORKING_SCHEMA      : Working_schema
  OPEN_CONTENTS       : set of Open_contents

Message_queue ::
  QUEUE_OBJECT : Object_designator
  MESSAGES     : seq of Message

Workstation ::
  WORKSTATION_OBJECT : Object_designator
  AUDIT_CRITERIA     : set of Selection_criterion

```

Instances of the PCTE concepts are called *entities*; they are referred to by the names of the underlying concepts. The state comprises the entities of a PCTE installation that endure from one operation call to another. The effect of an operation call is to modify the state, or to return values derived from the state (and any parameters), or both.

The system time is the date and time of day at any instant, as given by some system clock. For the format of the time see 23.1.1.5. The *current time* for an operation is a value of the system time at some moment between the start and end of the operation.

The object base is a set of objects identified by object designators (see 8.2.1).

A working schema is associated with a process (see clause 13) and consists of a set of types in working schema, derived from a sequence of SDSs. The types in working schema in the working schema of the calling process are called *visible types*. For the creation of a working schema for a process see 13.2.12.

The initial value of the state consists of the following objects:

- at least one workstation, at least one device managed by that workstation, at least one volume mounted on that device, and at least one process running on that workstation (see 18.1.2, 11.1.3, 11.1.1, and 13.1.5);
- the administration replica set, the common root, and the administrative objects (see 17.1.4 and 9.1.2);

- at least one user (see 19.1.1);
- at least the schema definition sets system, metasds, discretionary\_security, mandatory\_security (if implemented), and accounting (if implemented) (see 10.1);
- the predefined user group ALL\_USERS, and the predefined program groups PCTE\_AUDIT, PCTE\_REPLICATION, PCTE\_EXECUTION, PCTE\_SECURITY, PCTE\_HISTORY, PCTE\_CONFIGURATION, and PCTE\_SCHEMA\_UPDATE (see 19.1.1).

NOTE - It is intended that the system time should be as near as possible the same throughout a PCTE installation.

## 8.2 The object base

### 8.2.1 Objects

Object ::  
 OBJECT\_TYPE : Object\_type\_nominator  
 ATTRIBUTES : **set of** Attribute  
 LINKS : **set of** Link  
 DIRECT\_COMPONENTS : **set of** Object  
 PREFERRED\_LINK\_TYPE : [ Link\_type\_nominator ]  
 PREFERRED\_LINK\_KEY : [ Text ]  
 CONTENTS : [ Contents ]

Object\_designator :: Token

Object\_designators = **set of** Object\_designator

Contents = Structured\_contents | Unstructured\_contents

Structured\_contents = Accounting\_log | Audit\_file

Unstructured\_contents = File | Pipe | Device

Object\_scope = ATOMIC | COMPOSITE

The object type constrains the properties of the object (see 8.3.1).

No two attributes of an object have the same attribute type. There is a basic set of attributes which all objects have; it is defined in 9.1.1.

The preferred link type and preferred link key, if present, are used as defaults in the identification of a link of the object (see 8.2.3). The preferred link key has the syntax of a key (see 23.1.2.7).

Every direct component of an object is the destination of a composition link of the object, and vice versa.

An *outer object* of an object A is an object of which A is a component.

The *atomic object* associated with an object comprises the links, attributes, preferred link type, preferred link key, and contents of the object. The *atoms* of an object are the atomic objects associated with the object and all its components.

A *component* of an object is a direct component of the object or of a component of the object. An object which is a component of each of two distinct objects, neither of which is a component of the other, is called a *shared component* of those two objects.

An *internal link* of an object is a link of the object or of one of its components for which the destination is either a component of the object or the object itself. An *external link* of an object is a direct or indirect outgoing link of the object which is not an internal link of the object. An object is called the *origin* of each of its links.

An object is specified by an object designator, or by a specialization of object designator defined as follows: if "X" is an object type (that is, it is a descendant of "system-object", see 9.1.1) then 'X\_designator' (with capital initial) stands for 'Object\_designator' with the condition that the value must designate an object of type "X" or a descendant of "X". For the mapping of object designators to the language bindings, see 23.1.2.2.

An object scope is used to indicate whether the effect of an operation applies to an object (COMPOSITE) or to the atomic object of the object (ATOMIC).

#### NOTES

1 An object can be a component of itself. Similarly two objects can be components of each other; in that case there are two distinct objects with the same atoms.

2 General operations are provided for handling unstructured contents (see clause 12) as a sequence of octets, the meaning of which is not further defined in this part of ISO/IEC 13719. Specific operations are provided for handling structured contents, which has a defined meaning in each case (see clauses 21 and 22).

3 When an object is created, so are all its attributes in the global schema. When a new attribute type is applied to the object's type in an SDS, effectively all objects of that type and its descendants gain a new attribute with its initial value. If the application of an attribute type to that object type is removed from all SDSs, the attribute remains on each object of that type until deleted by OBJECT\_DELETE\_ATTRIBUTE.

## 8.2.2 Attributes

Attribute ::

ATTRIBUTE\_TYPE : Attribute\_type\_nominator

ATTRIBUTE\_VALUE : Attribute\_value

Attribute\_value = Integer | Natural | Boolean | Time | Float | String

Attribute\_designator :: Token

Attribute\_designators = **set of** Attribute\_designator

Attribute\_selection = Attribute\_type\_nominators | VISIBLE\_ATTRIBUTE\_TYPES

Attribute\_assignments = **map** Attribute\_designator **to** Attribute\_value

String = **seq of** Octet

The value of an enumeration attribute is represented by its position within the enumeration value type (see 8.3.2).

An attribute is specified as follows:

- for an attribute of an object: an object designator which specifies the object, and an attribute designator which specifies the attribute relative to the object;
- for an attribute of a link: an object designator which specifies the origin of the link, a link designator which specifies the link relative to the object (see 8.2.3), and an attribute designator which specifies the attribute relative to the link.



## NOTES

1 Each attribute in the object base is a key or non-key attribute of a link in the object base or a direct attribute of an object in the object base.

2 An implementation may impose constraints on the values of attributes (see clause 24). An attribute may take any value of its value type within those constraints; for example, a string attribute may take any string value up to the maximum allowed length, whatever its present value may be.

3 For the types Integer, Natural, Boolean, Time, Float, and String see 23.1.1.

## 8.2.3 Links

Link ::

|                    |                           |
|--------------------|---------------------------|
| LINK_TYPE          | : Link_type_nominator     |
| DESTINATION        | : [ Object_designator ]   |
| KEY_ATTRIBUTES     | : <b>seq of</b> Attribute |
| NON_KEY_ATTRIBUTES | : <b>set of</b> Attribute |
| REVERSE            | : [ Link_designator ]     |

Link\_designator :: Token

Actual\_key = **seq1 of** (Text | Natural)

Link\_designators = **set of** Link\_designator

Link\_selection = Link\_type\_nominators | VISIBLE\_LINK\_TYPES | ALL\_LINK\_TYPES

Link\_descriptor = Object\_designator \* Link\_designator

Link\_descriptors = **set of** Link\_descriptor

Link\_set\_descriptor = Object\_designator \* Link\_designators

Link\_set\_descriptors = **set of** Link\_set\_descriptor

Link\_scope = INTERNAL\_LINKS | EXTERNAL\_LINKS | ALL\_LINKS

The key attributes and the non-key attributes are together called the *attributes of the link*. No two attributes of a link have the same attribute type.

Two distinct links of the same type from the same object must have different key attributes (i.e. the two sequences of key attribute values must be different).

The reverse link of the reverse link of a link is that link.

A link is said to be *from* its origin and *to* its destination.

A *series of links* from object A to object B is a sequence of 1 or more links L1, L2, ..., Ln such that A is the origin of L1, B is the destination of Ln, and otherwise the destination of each link is the origin of the next in sequence.

A link is specified by an object designator which specifies the origin of the link and a link designator which specifies the link relative to the object. For the mapping of link designators to the language bindings, see 23.1.2.4.

## NOTES

1 Each link in the object base is a link of exactly one object in the object base; i.e. each link has exactly one origin.

2 When a link is created, so are all its attributes in the global schema. When a new attribute type is applied to the link's type in an SDS, effectively all links of that type gain a new attribute with its initial value. If the application of an attribute type to that link type is removed from all SDSs, the attribute remains on each link of that type until deleted by LINK\_DELETE\_ATTRIBUTE.



### 8.3 Types

Type = Object\_type | Attribute\_type | Link\_type | Enumeral\_type

Type\_nominator = Object\_type\_nominator | Attribute\_type\_nominator | Link\_type\_nominator | Enumeral\_type\_nominator

Object\_type\_nominator :: Token

Attribute\_type\_nominator :: Token

Link\_type\_nominator :: Token

Enumeral\_type\_nominator :: Token

Type\_nominators = **set of** Type\_nominator

Object\_type\_nominators = **set of** Object\_type\_nominator

Attribute\_type\_nominators = **set of** Attribute\_type\_nominator

Link\_type\_nominators = **set of** Link\_type\_nominator

Type\_kind = OBJECT\_TYPE | ATTRIBUTE\_TYPE | LINK\_TYPE | ENUMERAL\_TYPE

A type is a template defining common basic properties of a set of instances. The *instances* of a type are those whose type nominator identifies that type.

A type is specified by a type nominator, which may be specialized to an object type nominator, an attribute type nominator, a link type nominator, or an enumeral type nominator. A type nominator may be further specialized as follows: if "X" is an object type, attribute type, link type, or enumeral type then 'X\_type\_nominator' stands for 'Object\_type\_nominator' etc. with the condition that the value must designate type "X" or a descendant of "X". For the mapping of type nominators to language bindings see 23.1.2.5 and 23.1.2.

#### 8.3.1 Object types

Object\_type ::

TYPE\_NOMINATOR : Object\_type\_nominator

CONTENTS\_TYPE : [ Contents\_type ]

PARENT\_TYPES : Object\_type\_nominators

CHILD\_TYPES : Object\_type\_nominators

**represented by** object\_type

Contents\_type = FILE\_TYPE | PIPE\_TYPE | DEVICE\_TYPE | AUDIT\_FILE\_TYPE |  
ACCOUNTING\_LOG\_TYPE

The contents type, if present, specifies the type of contents of instances of the object type. If no contents type is supplied, instances of the object type have no contents.

The parent types define inheritance rules governing the properties of object types in working schema (see 8.5.1). The parent types of an object type, their parent types, and so on, excluding the object type itself, are called the *ancestor types* of the object type.

The child types are the object types which have this object type as parent type. The child types of an object type, their child types, and so on, excluding the object type itself, are called the *descendant types* of the object type.

The parent/child relation between object types forms a directed acyclic graph, with the object type "object" (see 9.1.1) as the root.

### 8.3.2 Attribute types

Attribute\_type ::  
 TYPE\_NOMINATOR : Attribute\_type\_nominator  
 VALUE\_TYPE\_IDENTIFIER : Value\_type\_identifier  
 INITIAL\_VALUE : [ Attribute\_value ]  
 DUPLICATION : Duplication  
**represented by** attribute\_type

Value\_type\_identifier = INTEGER | NATURAL | BOOLEAN | TIME | FLOAT | STRING |  
 Enumeration\_value\_type\_identifier

Enumeration\_value\_type\_identifier = **seq1 of** Enumeral\_type\_nominator

Duplication = DUPLICATED | NON\_DUPPLICATED

The value type identifier identifies the *value type* of the instances of the attribute type, i.e. the datatype of their possible attribute values (see table 1). See 23.1.1 for the mapping of values of integers, naturals, Booleans, times, floats, and strings. An enumeration value type identifier is a non-empty sequence of enumerals.

The initial value, which is a value of the value type, is the initial value of any attribute of this attribute type after creation and before any value has been assigned to it. If no initial value is supplied, the default initial value for the value type is used (see table 1).

If the duplication is DUPLICATED, then every instance of the attribute type is a *duplicable* attribute, i.e. the value of the attribute is copied whenever an object or link with the attribute is copied; if it is NON\_DUPPLICATED then every instance is a *nonduplicable* attribute, i.e. the value of the copy of the attribute reverts to the initial value.

Table 1 - Value types

| Value type identifier             | Value type    | Default initial value                                       |
|-----------------------------------|---------------|-------------------------------------------------------------|
| INTEGER                           | Integer       | 0                                                           |
| NATURAL                           | Natural       | 0                                                           |
| BOOLEAN                           | Boolean       | <b>false</b>                                                |
| TIME                              | Time          | 1980-01-01T00:00:00Z                                        |
| FLOAT                             | Float         | 0.0                                                         |
| STRING                            | String        | "" (empty string)                                           |
| Enumeration value type identifier | Enumeral type | 1st enumerual type of the enumeration value type identifier |

### 8.3.3 Link types

Link\_type ::  
 TYPE\_NOMINATOR : Link\_type\_nominator  
 CATEGORY : Category  
 LOWER\_BOUND, UPPER\_BOUND : [ Natural ]  
 EXCLUSIVENESS : Exclusiveness  
 STABILITY : Stability  
 DUPLICATION : Duplication  
 KEY\_ATTRIBUTE\_TYPES : Key\_types  
 REVERSE\_LINK\_TYPE : [ Link\_type\_nominator ]  
**represented by** link\_type

Key\_types = **seq of** Attribute\_type\_nominators

Category = COMPOSITION | EXISTENCE | REFERENCE | DESIGNATION | IMPLICIT

Categories = **set of** Category

Exclusiveness = SHARABLE | EXCLUSIVE

Stability = ATOMIC\_STABLE | COMPOSITE\_STABLE | NON\_STABLE

All instances of a link type have the category, exclusiveness, stability, and duplication of the link type.

The lower bound of a link type defines the number below which the number of links of that link type from any instance of an object type with that link type cannot be reduced. If absent, the lower bound is taken as 0. The lower bound is only checked when an attempt is made to delete a link, so that on creation of an object the number of links of a type may be less than the lower bound for that type.

The upper bound of a link type is an optional natural defining the maximal number of links of that link type from any instance of an object type with that link type. If present, it must be greater than 0 and not less than the lower bound. If absent, there is no upper bound.

A link type is said to be *of cardinality one* if its upper bound is 1. A link type of cardinality one has an empty sequence of key attribute types.

A link type is said to be of *cardinality many* if it is not of cardinality one. A link type of cardinality many has a non-empty sequence of key attribute types.

The sequence of key attribute types defines the attribute types of the sequence of key attributes of an instance of the link type. It does not contain any repeated attribute type nominators. A key attribute has value type Natural or String.

The optional reverse link type is the link type which reverses the link type, i.e. whenever a link of this link type exists from object A to object B, a link of the reverse type exists from object B to object A, and vice versa. The reverse link type is not allowed if the category is DESIGNATION, and must be present otherwise.

The term *complementary* is used of pairs of links, each having the other's origin as its destination, which are not reverses of each other.

All link types of category IMPLICIT and cardinality many have lower bound 0, no upper bound, and a single key attribute of the predefined attribute type "system\_key". The values of "system\_key" attributes are implementation-dependent: each such key value is different from the value of every other "system\_key" attribute of a link of the same link type from the same object.

All link types of category EXISTENCE and cardinality many have lower bound 0.

The category identifies certain *properties* of instances of the link type, as follows:

- *relevance to the origin.* For a link with this property:
  - . The link may be created and deleted explicitly.
  - . APPEND\_LINKS discretionary access right to the origin is required in order to create the link, and WRITE\_LINKS discretionary access right to the origin is required in order to delete the link.
  - . The link cannot be created or deleted if its origin is a stable object.
  - . The creation and deletion of the link assign the current system time to the last modification time of the origin.
  - . The link may have non-key attributes.

For a link without the relevance to the origin property:

- . The link may only be created and deleted implicitly, i.e. as the reverse of a link with the relevance to the origin property.
  - . APPEND\_IMPLICIT discretionary access right to the origin is required in order to create the link, and WRITE\_IMPLICIT discretionary access right to the origin is required in order to delete the link.
  - . The link can be created or deleted even if its origin is a stable object.
  - . The creation and deletion of the link have no effect on the last modification time of the origin.
  - . The link may not have non-key attributes.
- *referential integrity.* For a link with this property:
    - . If the link exists then so does its destination, i.e. the existence of the link prevents the deletion of its destination.

- . The link always has a reverse link with the referential integrity property.
- *existence property*. For a link with this property:
  - . An object can be created as destination of the link.
  - . The deletion of the link can imply the deletion of its destination.
- *composition property*. For a link with this property:
  - . The destination of the link is a component of its origin.

The categories are defined in terms of these properties as follows:

- COMPOSITION: relevance to the origin, referential integrity, existence property, composition property. Links with this category are called *composition links*.
- EXISTENCE: relevance to the origin, referential integrity, existence property. Links with this category are called *existence links*.
- REFERENCE: relevance to the origin, referential integrity. Links with this category are called *reference links*.
- IMPLICIT: referential integrity. Links with this category are called *implicit links*.
- DESIGNATION: relevance to the origin. Links with this category are called *designation links*.

If the stability of a link type is ATOMIC\_STABLE, each instance of the link type is an *atomically stabilizing* link, i.e. the destination of the link (excluding its components other than itself) cannot be modified or deleted.

If the stability of a link type is COMPOSITE\_STABLE, each instance of the link type is a *compositely stabilizing* link, i.e. the destination of the link (including its components) cannot be modified or deleted.

If the stability of a link type is NON\_STABLE, each instance of the link type is a *nonstabilizing* link, i.e. the existence of the link does not prevent the modification or deletion of its destination or its components.

Modification of an object is defined in 9.1.1. A *stable* object is the destination of an atomically or compositely stabilizing link, or a component of the destination of a compositely stabilizing link.

Exclusiveness applies only to composition link types. If it is EXCLUSIVE, each instance of the link type is an *exclusive* composition link, i.e. no other composition link can share the same destination. If it is SHARABLE, each instance of the link type is a *sharable* composition link, i.e. other composition links can share the same destination.

If duplication is DUPLICATED, each instance is a *duplicable* link, i.e. the link is copied whenever its origin is copied; if it is NON\_DUPLICATED, each instance is a *nonduplicable* link, i.e. a copy of the object has no copy of the link. An implicit link cannot be duplicable.

A component of an object is a *duplicable component* if it is the destination of at least one duplicable internal composition link whose origin is either the object or a duplicable component of the object.

A link type of category IMPLICIT or DESIGNATION must be nonstabilizing.

The following relations hold between properties of a link type and of its reverse link type:

- if one link type has category IMPLICIT, then the other does not;
- if one link type has the existence property (i.e. has category EXISTENCE or COMPOSITION) then the other does not;
- if one link type has stability ATOMIC\_STABLE or COMPOSITE\_STABLE then the other has category IMPLICIT.

A link type of category DESIGNATION cannot have a reverse link type.

Links of the following types are termed *usage designation links*, because they are not checked by the normal security rules: "running\_process", "in\_working\_schema\_of", "consumer\_process", "user\_identity\_of", "adopted\_user\_group\_of", "reserved\_by", "locked\_by", "lock", "opened\_by", "mounted\_on", and "listened\_to". Usage designation links have the following properties:

- creation or deletion implies only a bitwise write access on the origin object from a mandatory security point of view (see 20.1.8.2);
- creation or deletion requires one unspecified discretionary access permission on the origin object;
- creation or deletion is possible for an object on a read-only volume;
- creation or deletion is possible for a copy object as origin;
- creation or deletion does not require the establishment of locks on the links;
- they are not copied by REPLICATED\_OBJECT\_DUPLICATE;
- they can be implicitly deleted by network failure and workstation closedown;
- creation or deletion does not change the last modification time of the origin object.

Links of the following types are termed *service designation links*, because they indicate that the destination provides a service to the origin (usually a process): "executed\_on", "sds\_in\_working\_schema", "consumer\_identity", "user\_identity", "adopted\_user\_group", "reserved\_message\_queue", "open\_object", "process\_waiting\_for", "referenced\_object", "adoptable\_user\_group", "mounted\_volume", "is\_listener", "notifier", and "executed\_static\_context". Service designation links have the following properties:

- creation or deletion does not require the establishment of locks on the links;
- they are implicitly deleted by workstation failure;
- for navigation along these links to replicated objects, replication redirection applies to the state of the object base at the time the link was created rather than when it is navigated through.

#### NOTES

- 1 The properties of links of various categories are summarized in table 2.

Table 2 - Properties of link categories

| Property              | Composition links | Existence links | Reference links | Implicit links | Designation links |
|-----------------------|-------------------|-----------------|-----------------|----------------|-------------------|
| relevance to origin   | yes               | yes             | yes             | no             | yes               |
| referential integrity | yes               | yes             | yes             | yes            | no                |
| existence             | yes               | yes             | no              | no             | no                |
| composition           | yes               | no              | no              | no             | no                |
| atomic stability      | optional          | optional        | optional        | no             | no                |
| composite stability   | optional          | optional        | optional        | no             | no                |
| exclusiveness         | optional          | no              | no              | no             | no                |
| duplication           | optional          | optional        | optional        | no             | optional          |
| has a reverse link    | yes               | yes             | yes             | yes            | no                |

2 The reason why the lower bound of an existence link is 0 is that if there existed an existence link type L with a lower bound of 2, for example, and an object X had two outgoing links of type L, it would be impossible to delete either link directly using LINK\_DELETE. Indirect deletion of these links by deletion of object X would also be impossible because X would have outgoing existence links. This means that the destinations of these links could never be deleted. This would be an undesirable situation. The same problem does not exist with composition links because a composite object can be deleted in a single operation, OBJECT-DELETE.

### 8.3.4 Enumeral types

Enumeral\_type ::  
 TYPE\_NOMINATOR : Enumeral\_type\_nominator  
**represented by** enumeral\_type

An enumeral type is used as a possible value of an enumeration attribute. It has no instances.

## 8.4 Types in SDS

Type\_in\_sds = Object\_type\_in\_sds | Attribute\_type\_in\_sds | Link\_type\_in\_sds |  
 Enumeral\_type\_in\_sds

Type\_in\_sds\_common\_part ::  
 ASSOCIATED\_TYPE : Type\_nominator  
 LOCAL\_SDS : Object\_designator  
 LOCAL\_NAME : [ Name ]

Type\_nominator\_in\_sds = Object\_type\_nominator\_in\_sds | Attribute\_type\_nominator\_in\_sds |  
 Link\_type\_nominator\_in\_sds | Enumeral\_type\_nominator\_in\_sds

Object\_type\_nominator\_in\_sds :: Token

Attribute\_type\_nominator\_in\_sds :: Token

Link\_type\_nominator\_in\_sds :: Token

Enumeral\_type\_nominator\_in\_sds :: Token

Type\_nominators\_in\_sds = **set of** Type\_nominator\_in\_sds

Object\_type\_nominators\_in\_sds = **set of** Object\_type\_nominator\_in\_sds  
 Attribute\_type\_nominators\_in\_sds = **set of** Attribute\_type\_nominator\_in\_sds  
 Link\_type\_nominators\_in\_sds = **set of** Link\_type\_nominator\_in\_sds  
 Enumeral\_type\_nominators\_in\_sds = **set of** Enumeral\_type\_nominator\_in\_sds  
 Definition\_mode\_value = CREATE\_MODE | DELETE\_MODE | READ\_MODE | WRITE\_MODE | NAVIGATE\_MODE  
 Definition\_mode\_values = **set of** Definition\_mode\_value  
 Definition\_modes ::  
     USAGE\_MODE : Definition\_mode\_values  
     EXPORT\_MODE : Definition\_mode\_values  
     MAXIMUM\_USAGE\_MODE : Definition\_mode\_values

A type in SDS (plural 'types in SDS') is a template defining a set of properties which apply to all instances of its type, in addition to the basic properties of that type. A type in SDS is *associated with* one type; a type is *associated with* one or more types in SDS.

A *schema definition set* (or *SDS*) is an object of type "sds" (see 10.1.1), and is specified by an object designator. A type in SDS *belongs to*, or is *in*, a particular SDS, called its local SDS.

The local name identifies the type in SDS, and hence the associated type, uniquely within the local SDS. The *complete name* of a type in SDS is the name of the SDS, followed by a hyphen '-', followed by the local name of the type in SDS.

The definition modes specify restrictions on the usage of the type in SDS. The usage mode specifies the permitted kinds of access to instances of the type in SDS by a process which has adopted its local SDS in its working schema. The export mode specifies the maximum usage mode of the copy of the type in SDS which is created when the type in SDS is exported to another SDS; it is a subset of the usage mode. The maximum usage mode specifies which definition mode values can be included in the usage mode and export mode; it is set on creation of the type in SDS and cannot be changed. The definition modes of a link and of its reverse must be the same. Enumeral types in SDS do not have definition modes.

The accesses controlled by definition modes are as follows.

- READ\_MODE controls reading from attributes by the operations OBJECT\_GET\_ATTRIBUTE, OBJECT\_GET\_SEVERAL\_ATTRIBUTES, LINK\_GET\_ATTRIBUTE, and LINK\_GET\_SEVERAL\_ATTRIBUTES.
- WRITE\_MODE controls writing to attributes by the operations OBJECT\_SET\_ATTRIBUTE, OBJECT\_SET\_SEVERAL\_ATTRIBUTES, LINK\_SET\_ATTRIBUTE, LINK\_SET\_SEVERAL\_ATTRIBUTES, OBJECT\_RESET\_ATTRIBUTE, and LINK\_RESET\_ATTRIBUTE.
- CREATE\_MODE controls creation of objects and links by the operations OBJECT\_CREATE, OBJECT\_COPY, OBJECT\_CONVERT, VERSION\_REVISE, VERSION\_SNAPSHOT, DEVICE\_CREATE, and PROCESS\_CREATE; and creation of links by the operations LINK\_CREATE and LINK\_REPLACE.
- DELETE\_MODE controls deletion of objects and links by the operation OBJECT\_DELETE and deletion of links by the operations LINK\_DELETE and LINK\_REPLACE.



- NAVIGATE\_MODE controls the use of link references in pathnames in the evaluation of object references (see 23.1.2.2).

Types in SDS are specialized to object types in SDS, attribute types in SDS, link types in SDS, and enumerals in SDS; the associated types are object types, attribute types, link types, and enumerals respectively.

A type in SDS is specified by a type nominator in SDS, which may be specialized to an object type nominator in SDS, an attribute type nominator in SDS, a link type nominator in SDS, or an enumerals type nominator in SDS. A type nominator in SDS may be further specialized as follows: if "X" is an object type, attribute type, link type, or enumerals type then 'X\_type\_nominator\_in\_sds' stands for 'Object\_type\_nominator\_in\_sds' etc. with the condition that the value must designate a type in SDS associated with type "X" or a descendant of "X". For the mapping of type nominators in SDS to language bindings see 23.1.2.5.

NOTE - The properties of a type and of an associated type in SDS can be specified by means of the Data Definition Language (see annex B).

#### 8.4.1 Object types in SDS

```
Object_type_in_sds :: Type_in_sds_common_part &&
    DIRECT_ATTRIBUTE_TYPES_IN_SDS      : Attribute_type_nominators_in_sds
    DIRECT_OUTGOING_LINK_TYPES_IN_SDS   : Link_type_nominators_in_sds
    DIRECT_COMPONENT_TYPES_IN_SDS      : Object_type_nominators_in_sds
    DEFINITION_MODES                    : Definition_modes
represented by object_type_in_sds
```

The only allowed definition mode value for an object type in SDS is CREATE\_MODE.

The direct attribute types in SDS must be in the same SDS as the object type in SDS. They participate in the definition of the visible attribute types of object types in working schema; see 8.5.1.

The direct outgoing link types in SDS must be in the same SDS as the object type in SDS. They participate in the definition of the visible link types of object types in working schema; see 8.5.1. The object type in SDS is called the *origin object type in SDS* of each of the direct outgoing link types in SDS.

The direct component types in SDS must be in the same SDS as the object type in SDS. They participate in the definition of the direct component types of object types in working schema; see 8.5.1.

#### 8.4.2 Attribute types in SDS

```
Attribute_type_in_sds :: Type_in_sds_common_part &&
    DEFINITION_MODES : Definition_modes
represented by attribute_type_in_sds
```

The only allowed definition mode values for an attribute type in SDS are READ\_MODE and WRITE\_MODE.

### 8.4.3 Link types in SDS

```

Link_type_in_sds :: Type_in_sds_common_part &&
  DESTINATION_OBJECT_TYPES_IN_SDS      : Object_type_nominators_in_sds
  NON_KEY_ATTRIBUTE_TYPES_IN_SDS        : Attribute_type_nominators_in_sds
  DEFINITION_MODES                      : Definition_modes
represented by link_type_in_sds

```

The only allowed definition mode values for a link type in SDS are CREATE\_MODE, DELETE\_MODE, and NAVIGATE\_MODE.

The destination object types in SDS must be in the same SDS as the link type in SDS. They participate in the definition of the destination object types of link types in working schema; see 8.5.3.

The non-key attribute types in SDS must be in the same SDS as the link type in SDS. They participate in the definition of the visible attribute types of link types in working schema; see 8.5.3.

### 8.4.4 Enumeral types in SDS

```

Enumeral_type_in_sds :: Type_in_sds_common_part &&
  IMAGE      : Text
represented by enumeral_type_in_sds

```

An enumeral type in SDS associates with the enumeral type a string called its image.

## 8.5 Types in working schema

```

Type_in_working_schema = Object_type_in_working_schema |
  Attribute_type_in_working_schema | Link_type_in_working_schema |
  Enumeral_type_in_working_schema

Type_in_working_schema_common_part ::
  ASSOCIATED_TYPE      : Type_nominator
  CONSTITUENT_TYPES_IN_SDS : seq of (Composite_name * Type_nominator_in_sds)

Composite_name ::
  SDS_NAME      : Name
  LOCAL_NAME    : [ Name ]

```

A type in working schema is a template defining common properties for a set of instances of its type. The properties of a type in working schema are derived from the properties of one or more types in SDS (see 8.5.1 to 8.5.4). Types in working schema occur in working schemas, see 8.1. For the construction of working schemas, see 13.2.12.

The constituent types in SDS of a type in working schema must all have the same associated type, which is the type *associated with* the type in working schema.

A type in working schema has several composite names, one for each constituent type in SDS. For each composite name, the SDS name is the name of the local SDS of the corresponding type in SDS, and the local name, if any, is the local name of the type in SDS in its local SDS.

Let C1 and C2 be composite names, and T1 and T2 be type nominators in SDS. Then for any two constituent types in SDS (C1, T1), (C2, T2) of a type in working schema, if the SDS name of

C1 precedes the SDS name of C2 in the SDS names of the working schema containing the type in working schema, then (C1, T1) precedes (C2, T2).

Types in working schema are specialized to object types in working schema, attribute types in working schema, link types in working schema, and enumeral types in working schema; their associated types are object types, attribute types, link types, and enumeral types respectively.

The value of a type in SDS cannot be changed while it is part of a type in working schema.

A type in working schema is specified by a type nominator, see 8.3.

### 8.5.1 Object types in working schema

```
Object_type_in_working_schema :: Type_in_working_schema_common_part &&
CHILD_TYPES                : Object_type_nominators
PARENT_TYPES                : Object_type_nominators
APPLIED_ATTRIBUTE_TYPES     : Attribute_type_nominators
APPLIED_LINK_TYPES          : Link_type_nominators
VISIBLE_ATTRIBUTE_TYPES     : Attribute_type_nominators
VISIBLE_LINK_TYPES          : Link_type_nominators
DIRECT_COMPONENT_TYPES     : Object_type_nominators
USAGE_MODES                 : Definition_mode_values
```

The set of constituent types in SDS of the child types is the union of the sets of child types of the constituent types in SDS of the type in working schema.

The set of constituent types in SDS of the parent types is the union of the sets of parent types of constituent types in SDS of the type in working schema.

The applied attribute types are the attribute types in working schema which have a constituent type in SDS of a direct attribute type in SDS of one of the constituent types in SDS of the object type in working schema.

The applied link types are the link types in working schema which have a constituent type in SDS of a direct outgoing link type in SDS of one of the constituent types in SDS of the object type in working schema.

The direct component types are the object types in working schema which have a constituent type in SDS of a direct component type in SDS of one of the constituent types in SDS of the object type in working schema.

The set of visible attribute types is the union of the set of applied attribute types and the sets of the visible attribute types of all the parent types.

The set of visible link types is the union of the set of applied link types and the sets of the visible link types of all the parent types.

The constituent types in SDS must be object types in SDS.

If the type of an object is not visible, the object is considered as an instance of any of its object type's ancestor types which is visible.

The set of usage modes is the union of the sets of definition modes of all constituent types in SDS of the object type in working schema.

### 8.5.2 Attribute types in working schema

Attribute\_type\_in\_working\_schema :: Type\_in\_working\_schema\_common\_part &&  
 USAGE\_MODES : set of Definition\_mode\_values

The constituent types in SDS must be attribute types in SDS.

The set of usage modes is the union of the sets of definition modes of all constituent types in SDS of the attribute type in working schema.

### 8.5.3 Link types in working schema

Link\_type\_in\_working\_schema :: Type\_in\_working\_schema\_common\_part &&  
 DESTINATION\_OBJECT\_TYPES : Object\_type\_nominators  
 VISIBLE\_DESTINATION\_OBJECT\_TYPES : Object\_type\_nominators  
 KEY\_ATTRIBUTE\_TYPES : Key\_types  
 APPLIED\_ATTRIBUTE\_TYPES : Attribute\_type\_nominators  
 REVERSE : [ Link\_type\_nominator ]  
 USAGE\_MODES : Definition\_mode\_values

The set of constituent types in SDS of the applied attribute types is the union of the sets of non-key attribute types of the constituent types in SDS of the link type in working schema.

The set of constituent types in SDS of the destination object types is the union of the sets of destination object types of the constituent types in SDS of the link type in working schema.

The constituent types in SDS must be link types in SDS.

The set of visible destination object types is the union of the set of destination object types and the set of visible descendants of the visible destination object types.

The sequence of key attribute types is the same as the sequence of key attribute types of the associated type.

The set of usage modes is the union of the sets of definition modes of all constituent types in SDS of the link type in working schema.

### 8.5.4 Enumeral types in working schema

Enumeral\_type\_in\_working\_schema :: Type\_in\_working\_schema\_common\_part &&  
 IMAGE : Text

The image of an enumeral type in working schema T1 is the image of the first of its types in SDS which has an image, unless another enumeral type in working schema T2 belonging to the same enumeration attribute type in working schema as T1 but with a lower position already has the same image, in which case T1 has no image.

The constituent types in SDS must be enumeral types in SDS.

## 8.6 Types in global schema

The *global schema* is the working schema constituted by all the SDSs of a PCTE installation; the order is irrelevant as it affects only the type names, which are of no concern here. A *type in global schema* is a type in working schema in the global schema; it follows that each type is associated with one type in global schema. The global schema is a notional working schema used

to state the following consistency rules applying to the whole object base; it is not necessarily the working schema of any process.

An object must be compatible with its associated object type in global schema, i.e.:

- The link types in global schema of the links of the object must be among the visible link types in global schema of the object type in global schema.
- The attribute types in global schema of the attributes of the object must be among the visible attribute types in global schema of the object type in global schema.
- The object types in global schema of the direct components of the object must be among the direct component object types in global schema of the object type in global schema.
- The preferred link type of the object, if present, must be one of the applied link types of the object type in global schema.
- The preferred link key of the object, if present, must have the same value types (String or Natural), in the same order, as the key attribute types of the preferred link type of the object.

A link must be compatible with its associated link type in global schema, i.e.:

- The object type in global schema of the destination of the link must be among the visible destination types of the link type in global schema.
- The key attributes of the link, if any, must have the same value types (String or Natural) in the same order as the key attribute types of the link type in global schema.
- The non-key attributes of the link must be among the applied attribute types of the link type in global schema.
- The link type in global schema of the reverse link, if any, must be the reverse of the link type in global schema.

## 8.7 Operations

### 8.7.1 Calling process

The operations defined in clauses 9 to 22 take effect when they are *executed* by a process (see 13.1.4). The process is known as the *calling process* of the operation. The effects on the state of the PCTE installation are global, i.e. can be observed by other processes. Results returned by operations which are designators are local to the calling process.

### 8.7.2 Direct and indirect effects

The effects of an operation on the state of the PCTE installation comprise direct effects and indirect effects. *Direct effects* are described in the relevant operation descriptions (including the error descriptions). *Indirect effects* are described elsewhere in clauses 9 to 22. The operations of clause 23 do not affect the state.

Indirect effects occur by means of *events*. Events are of several classes, described below. An operation may *raise* an event. Depending on the state of the PCTE installation, the raising of an event may result either in the effect of another operation being different to what it would otherwise be, or in some other change of state.

An operation has a finite non-zero duration, and an event that is raised during the operation may have an effect on that operation.

In general, the raising of an event is not explicitly described in the operation that raises the event, but instead in the part of the interface definition that may be affected by the event. It must nevertheless be understood that the description of an operation may need to be implicitly extended by the description of the raising of events. The processing of an event takes place asynchronously.

There are several different classes of event, as described below.

- *Access event.* This is described in clause 15. Access events are raised by operations which perform certain kinds of access to objects. If an appropriate notifier has been established, then the raising of the event causes an appropriate message to be sent.
- *Lock release event.* A lock release event occurs when a lock is released (see 16.1.8). If some other operation is waiting to acquire a lock on the same resource, then that operation may proceed. If there is more than one such operation then which operation acquires the lock is implementation-dependent. There is no further description of this event in this part of ISO/IEC 13719.
- *Process timeout event.* This event is raised when the duration of an operation has exceeded the process timeout value for that process. When this event is raised, the error condition OPERATION\_HAS\_TIMED\_OUT holds for that operation, and the operation terminates with that error. This event is described in 13.1.4.
- *Process alarm event.* This event is raised when the time left until alarm has expired. When this event is raised, a message of message type WAKE is sent to the process and the process is resumed. This event and its effect are described in 13.1.4 and 13.2.6.
- *Interrupt operation event.* This is described in 13.2.4. This event is raised by PROCESS\_INTERRUPT\_OPERATION or by PROCESS\_TERMINATE. If the interrupted process is executing an operation or waiting for an event when this error is raised, then the error condition OPERATION\_IS\_INTERRUPTED holds for that operation and it terminates with that error.
- *Audit event.* These events are described in clause 21. Audit events are raised by operations which carry out object access, and for exploiting audit records, copying audit records, carrying out certain security operations, and at explicit user request (see 21.2.5). If the event type is selected and auditing is enabled (or the event type is always audited) then an audit record is written as described in 21.1.1.
- *Accounting event.* Accounting events are divided into start events and end events. These are raised as a result of certain operations, and if the resource is accountable and accounting is enabled then an accounting record is written (see 22.1.2).
- *Message queue event.* These events are described in 14.1. They are raised by the appearance in a message queue of a message, which may be caused by MESSAGE\_SEND\_WAIT, MESSAGE\_SEND\_NO\_WAIT, or QUEUE\_RESTORE. If there is a handler for the event then that handler is executed.
- *Resource availability event.* These events do not occur as a result of operations, but may occur at any moment. They model changes in the availability of hardware resources. The effect of a resource availability event on the directly affected objects is implementation-



dependent. The set of directly affected objects for an event is implementation-defined. The effect on access from other objects is defined as follows for the various resource availability events.

- *Volume failure*: an accessible volume becomes inaccessible. Attempted access to objects on the volume (including replicas in the case of an administration volume) fails with the error OBJECT\_IS\_INACCESSIBLE. Attempts to commit transactions which have started and which involve objects on the volume also fail.
  - *Device failure*: an accessible device becomes inaccessible. Attempted access to the file contents of the device fails. Volume failure is raised for any volume mounted on the device.
  - *Network failure*: an accessible workstation becomes inaccessible. There is no distinction between a workstation failing and a network failing so that communication with the workstation is lost. The inaccessible workstation ceases to be in its current network partition. Associated devices and volumes become inaccessible with consequences as above.
  - *Network repair*: a workstation joins a network partition. This has no immediate effect, but the workstation becomes accessible when the other conditions are met.
- *Process termination event*. This event is raised when a process is terminated by PROCESS\_TERMINATE, explicitly or implicitly or by normal or abnormal process termination. If a PROCESS\_WAIT\_FOR\_CHILD or PROCESS\_WAIT\_FOR\_ANY\_CHILD operation of the parent process is waiting for that process or any sibling process to terminate, then it may proceed. If more than one such operation exists (for different threads) then all may proceed.
- *Data available event*. This event is raised when data is written to a device contents, pipe contents, or message queue. If an operation is waiting on a CONTENTS\_READ on a pipe or device which is not non-blocking, and data is written to that pipe or device, then that operation may proceed. If an operation is waiting on a MESSAGE\_RECEIVE\_WAIT on a message queue, and a message is received of a type included in the set of types specified by that MESSAGE\_RECEIVE\_WAIT, then the operation may proceed. If more than one operation is waiting on that event, which operation receives the data and ceases to wait is implementation-dependent. AUDIT\_FILE\_READ and ACCOUNTING\_LOG\_READ do not wait for data to be available.
- *Data space available event*. This event is raised when space becomes available in a device contents, pipe contents, or message queue. If an operation is waiting on a CONTENTS\_WRITE on a pipe or device which is not non-blocking, and data is removed from that pipe or device, then that operation may proceed if sufficient space has been made available. If an operation is waiting on a MESSAGE\_SEND\_WAIT on a message queue and a message is removed from the queue then the operation may proceed if sufficient space has been made available. If an operation is waiting to write to the audit file or accounting log and the audit file or accounting log becomes available, then the operation proceeds. If more than one operation is waiting on that event for a contents or message queue, which operation writes or sends the data and ceases to wait is implementation-dependent.
- *Security attribute change*. This event is raised by OBJECT\_SET\_CONFIDENTIALITY\_LABEL, OBJECT\_SET\_INTEGRITY\_LABEL, or OBJECT\_SET\_ACL\_ENTRY, or

changes to labels as a result of floating. If a CONTENTS\_READ, CONTENTS\_WRITE, MESSAGE\_RECEIVE\_WAIT or MESSAGE\_SEND\_WAIT operation is waiting and a security attribute changes such that the process no longer has permission to access the contents or message queue object in the required mode, then the operation ceases to wait and terminates in an appropriate mandatory or discretionary access mode error (see C.4).

### 8.7.3 Errors

Execution of an operation may *terminate* after carrying out the *normal behaviour* as described in the main subclause, or may terminate in an *error*.

The list of errors in the subclause **Errors** of each abstract operation definition defines the set of *error conditions* which apply to that operation. Other error conditions, which apply to several operations, are defined in clause 23. The error conditions OPERATION\_HAS\_TIMED\_OUT and OPERATION\_IS\_INTERRUPTED (see 8.7.2) and SECURITY\_POLICY\_WOULD\_BE\_VIOLATED (see 20.1.8) apply to all operations. If an operation terminates in an error then the associated error condition holds. If any error condition holds then the operation terminates; if none of the error conditions hold, the normal behaviour occurs.

Error conditions are distinguished for the purpose of helping tool writers. Language bindings may add further error conditions. An implementation may not add further error conditions, except as specified in this part of ISO/IEC 13719.

No precedences are defined in this part of ISO/IEC 13719 between error conditions which hold simultaneously. Implementations which aim for high security must define such a precedence so as to address the problem of covert channels.

Error conditions arising from type mismatches between actual and formal parameters of operations are not explicitly defined in the abstract operation definitions. Language bindings may need to make these error conditions explicit, depending on the strength of type-checking provided by the language. This does not apply to the following cases:

- a check by an operation that an object belongs to a particular subset of instances of a type, e.g. that a security group is not a subgroup, or that an attribute is not applied to a specified object;
- a check by an operation where a specialization of Object\_designator is specified and an object reference is supplied (see 23.1.2.2).

If an operation terminates with an error condition, then the operation may have acquired some locks. The locks acquired are implementation-dependent, but in no case may a lock be acquired on a resource (object or link) which is stronger than the lock that would be acquired on that resource if the normal behaviour had occurred. Their duration is determined in the same way as for other locks. No other state changes occur, except that possibly auditing and accounting records are created.

Calls to operations which are part of optional modules which are not implemented by an implementation return with no error and no effect.



### 8.7.4 Operation serializability

In general, operations are serializable with all other concurrent operations. An operation may be considered to be composed of one or more atomic *actions* which change the state of the PCTE installation. That a set of operations is serializable means that for each operation a single point in time can be determined, lying between the actual time the operation is called and the time of the return from the operation, where all the actions of the operation can be deemed to take place and without any different effect on the state of the PCTE installation to that which actually occurs. This point of time is called the *nominal serialization point*. The following specific exceptions to serializability apply.

- If an operation enters a waiting state then the actions before the operation waits constitute one operation for purposes of serialization and the actions after leaving the waiting state until entering a further waiting state or the end of the operation constitute another operation.
- The values of the "last\_access\_time", "last\_modification\_time", "last\_change\_time", "last\_composite\_access\_time", "last\_composite\_modification\_time", and "last\_composite\_change\_time" attributes are updated within an operation to a point of time between the start and end of the operation and not necessarily to any nominal serialization point. The time values set on different objects by a single operation are not necessarily the same.
- If PROCESS\_INTERRUPT\_OPERATION is used on a process between the start of an operation and the nominal serialization point then the operation is interrupted; if it is used after the nominal serialization point then it has no effect.
- Serializability does not apply to PROCESS\_SUSPEND for the calling process; nor to WORKSTATION\_CHANGE\_CONNECTION with the parameter *force* set to **true** and any affected concurrent operations.
- PROCESS\_TERMINATE interrupts other ongoing operations (if any) in the same way as PROCESS\_INTERRUPT\_OPERATION.

#### NOTES

1 Serializability is often enforced by locking. However, this is not true for two or more operations running on behalf of the same activity or on behalf of competing unprotected activities. As an example of operation serializability, consider two concurrent invocations of OBJECT\_MOVE on the same object, moving it to two different volumes. The result should be that the entire object resides on one or the other volume, rather than some components residing on each volume according to the order in which they were moved.

2 Evaluation of parameters which are references counts as part of operation execution for serialization.

## 9 Object management

### 9.1 Object management concepts

#### 9.1.1 The basic type "object"

**sds** system:

volume\_identifier: (read) **non\_duplicated natural**;

locked\_link\_name: (read) **string**;

lock\_identifier: (read) **string**;

exact\_identifier: (read) **non\_duplicated string**;

```

number: natural;
name: string;
system_key: (read) natural;
object: with
attribute
    exact_identifier;
    volume_identifier;
    replicated_state: (read) non_duplicated enumeration (NORMAL, MASTER, COPY) :=
        NORMAL;
    last_access_time: (read) non_duplicated time;
    last_modification_time: (read) non_duplicated time;
    last_change_time: (read) non_duplicated time;
    last_composite_access_time: (read) non_duplicated time;
    last_composite_modif_time: (read) non_duplicated time;
    last_composite_change_time: (read) non_duplicated time;
    num_incoming_links: (read) non_duplicated natural;
    num_incoming_composition_links: (read) non_duplicated natural;
    num_incoming_existence_links: (read) non_duplicated natural;
    num_incoming_reference_links: (read) non_duplicated natural;
    num_incoming_stabilizing_links: (read) non_duplicated natural;
    num_outgoing_composition_links: (read) non_duplicated natural;
    num_outgoing_existence_links: (read) non_duplicated natural;
link
    predecessor: (navigate) non_duplicated composite stable existence link
        (predecessor_number: natural) to object reverse successor;
    successor: (navigate) implicit link (system_key) to object reverse predecessor;
    opened_by: (navigate) non_duplicated designation link (number) to process;
    locked_by: (navigate) non_duplicated designation link (lock_identifier) to activity
with attribute
    locked_link_name;
end locked_by;
end object;
end system;

```

"Object" is the common ancestor type of all objects in the object base.

The exact identifier uniquely identifies the object in the object bases of all PCTE installations. It is composed of a prefix and a suffix separated by ':' (colon). The prefix is the same for all objects created within a PCTE installation. The suffix uniquely identifies the object within the object base of a particular PCTE installation. The exact identifier of an object that has been deleted is never reassigned to an object created later.

The volume identifier identifies the volume on which the object resides, or, for a copy object, on which it is a replica. It uniquely identifies a volume within a PCTE installation.

The replicated state indicates whether the object is normal, a master or a copy (see 17.1). It is MASTER for the master of a replicated object, COPY for a copy of a replicated object, and NORMAL for a non-replicated object. This attribute can be changed only by the operations which manage replicated objects.

The last access time is the date and time of day of the last read access to the contents of the object. It is set to the system time when the object is created and by the following operations

(unless the object is on a read-only volume or is a component of an object on a read-only volume):

- QUEUE\_RESTORE (*queue, file*) for *file*;
- CONTENTS\_READ;
- AUDIT\_FILE\_READ;
- ACCOUNTING\_RECORD\_READ;
- CONTENTS\_COPY\_TO\_FOREIGN\_SYSTEM (*file, foreign\_system, foreign\_parameters, foreign\_name*) for *file*.

The last modification time is the date and time of day of the last *modification* to the object. It is set to the system time when the object is created and by the following operations:

- LINK\_CREATE and LINK\_REPLACE for the origin of the created link when the created link is not implicit;
- LINK\_DELETE and LINK\_REPLACE for the origin of the deleted link when the deleted link is not implicit;
- any operation which results in the creation or deletion of a link which is not implicit, except for usage designation links and "object\_on\_volume" links;
- OBJECT\_CONVERT;
- OBJECT\_SET\_ATTRIBUTE and OBJECT\_SET\_SEVERAL\_ATTRIBUTES;
- LINK\_SET\_ATTRIBUTE and LINK\_SET\_SEVERAL\_ATTRIBUTES, for the origins of the links;
- OBJECT\_SET\_PREFERENCE;
- QUEUE\_SAVE (*queue, file*) for *file*;
- CONTENTS\_WRITE and CONTENTS\_TRUNCATE;
- CONTENTS\_COPY\_FROM\_FOREIGN\_SYSTEM (*file, foreign\_system, foreign\_parameters, foreign\_name*) for *file*;
- any operation resulting in the creation of an audit record for the audit file;
- any operation resulting in the creation of an accounting record for the accounting log;

The last change time is the date and time of day of the last *change* to the object. It is set by any operation which sets the last modification time, and the following operations:

- creation of an implicit link;
- deletion of an implicit link;
- OBJECT\_MOVE for an object which has been moved to another volume;
- operations which change the discretionary access control lists of an object;
- operations which change the mandatory labels of an object;
- operations which change the mandatory label ranges of multi-level secure devices (see 20.1.5);

- PROCESS\_SET\_CONFIDENTIALITY\_LABEL (*process, label*) for *process*;
- PROCESS\_SET\_INTEGRITY\_LABEL (*process, label*) for *process*.

The last composite access time is the date and time of day of the last read access to the contents of the object or of any component of the object (but is not updated if the object or component is on a read-only volume).

The last composite modification time is the date and time of day of the last modification to the object or to any component of the object.

The last composite change time is the date and time of day of the last change made to the object or to any component of the object.

An operation which updates the last modification time of an object is said to *atomically modify* the object. An operation which updates the last composite modification time of an object is said to *compositely modify* the object.

The num (number of) incoming links is the number of non-designation links to the object (and is also the number of non-designation links of the object since every non-designation link has a reverse link).

The num (number of) incoming composition links is the number of composition links to the object.

The num (number of) incoming existence links is the number of existence links to the object.

The num (number of) incoming reference links is the number of reference links to the object.

The num (number of) incoming stabilizing links is the number of atomically stabilizing links to the object plus the number of compositely stabilizing links to the object and to its outer objects.

The num (number of) outgoing composition links is the number of composition links of the object.

The num (number of) outgoing existence links is the number of existence links of the object.

The destinations of the "predecessor" links are the immediate predecessor versions of the object; the destinations of the "successor" links are the immediate successor versions of the object. These are used in version control operations; see 9.4. The directed graph of versions created by these links must be acyclic.

The destination of the "opened\_by" link is the process that opened the object; see 12.1.

The destination of the "locked\_by" link is the activity that has locked the object or a link of the object; see 16.1.2.

There are also attributes defined in the security SDS representing the security properties of the object (see 19.1.2 and 20.1.1); and attributes defined in the accounting SDS representing the accounting properties of the object (see 22.1.1).

The attribute types "number", "name" and "system\_key" are predefined. "number" and "name" are used for numeric and string keys, respectively; names are non-empty. "system\_key" is the attribute type of system-assigned keys of implicit links (see 8.3.3).

#### NOTES

1 The prefix of the object exact identifier is intended to be unique among all PCTE installations, past, present, and future; but the administration of prefix assignment is outside the scope of this part of ISO/IEC 13719.

2 The last composite access, modification, and change time may be calculated when required as the most recent of the last access, modification, and change time respectively of the object and its components.

### 9.1.2 The common root

**sds** system:

common\_root: **child type of object with link**

archives: (**navigate**) **existence link to** archive\_directory **reverse** archives\_of;  
 execution\_sites: (**navigate**) **existence link to** execution\_site\_directory **reverse** execution\_sites\_of;  
 ground: (**protected**) **existence link to** common\_root;  
 replica\_sets: (**navigate**) **existence link to** replica\_set\_directory **reverse** replica\_sets\_of;  
 volumes: (**navigate**) **existence link to** volume\_directory **reverse** volumes\_of;

**end** common\_root;

**end** system;

The common root has an existence link to each of the *administrative objects* of the PCTE installation: the SDS directory (see 10.1.1), the volume directory (see 11.1.1), the archive directory (see 11.1.4), the replica set directory (see 17.1.1), the execution site directory, (see 18.1.1), the security group directory (see 19.1.1), the mandatory directory (see 20.1.1), and the accounting directory (see 22.1.1).

The "ground" link is to allow deletion of other existence links to the common root (see 9.2.2).

The common root and the administrative objects are predefined replicated objects (see 17.1.4).

### 9.1.3 Datatypes for object management

Type\_ancestry = EQUAL\_TYPE | ANCESTOR\_TYPE | DESCENDANT\_TYPE | UNRELATED\_TYPE

Version\_relation = ANCESTOR\_VSN | DESCENDANT\_VSN | SAME\_VSN | RELATED\_VSN | UNRELATED\_VSN

These datatypes are used as parameter and result types of operations in 9.3 and 9.4.

## 9.2 Link operations

### 9.2.1 LINK\_CREATE

```
LINK_CREATE (
  origin      : Object_designator,
  new_link    : Link_designator,
  dest        : Object_designator,
  reverse_key : [ Actual_key ]
)
```

LINK\_CREATE creates a new link *link* of *origin* as follows:

- the link type is as specified by the link designator *new\_link*;
- the destination is the object *dest*;
- the key attributes are as specified by the link designator *new\_link*;

- the non-key attributes are set to their initial values;
- the category, exclusiveness, stability, and duplication are set according to the link type.

If the type of *link* has a reverse link type, LINK\_CREATE also creates the reverse link *reverse\_link* of *link* and adds it to the links of *dest*:

- the link type of *reverse\_link* is the reverse of the link type of *link*;
- the destination of *reverse\_link* is *origin*;
- the category, exclusiveness, stability, and duplication of *reverse\_link* are set according to the link type.
- the non-key attributes of *reverse\_link* are set to their initial values;
- if the type of *reverse\_link* is of cardinality many and of category IMPLICIT then the key attribute of *reverse\_link* is set to a new system-generated key.

If the type of *reverse\_link* is of category COMPOSITION, REFERENCE or EXISTENCE, two cases arise:

- if *dest* has a preferred link type which is the link type of the reverse link, then the key attributes of *reverse\_link* are derived from *reverse\_key* and from the preferred link key of *dest*, if any, as defined in 23.1.2.7;
- if *dest* has no preferred link type or if the preferred link type of *dest* is not the link type of the reverse link, then the key of the reverse link is set to *reverse\_key*.

If *new\_link* is a composition link, then any security group that has OWNER granted or denied to *origin* has OWNER or CONTROL\_DISCRETIONARY granted or denied respectively to *dest*; similarly if *reverse\_link* is a composition link, then any security group that has OWNER granted or denied to *dest* has OWNER or CONTROL\_DISCRETIONARY granted or denied respectively to *origin*. This requires the process to have OWNER rights on *dest* or *origin* respectively. See 19.1.2 for details.

Write locks of the default mode are obtained on the new links. A read lock of the default mode is obtained on *origin* if the interpretation of *new\_link* implies the evaluation of any '+' or '++' key attribute values (see 23.1.2.7). A read lock of the default mode is obtained on *dest* if the interpretation of *reverse\_key* implies the evaluation of any '+' or '++' key attribute values.

A write lock of the default mode is obtained on *dest* and each of its components if the OWNER discretionary access right is granted or denied for one or more groups to *origin*, and different OWNER discretionary access rights exist for one or more of those same groups to *dest*.

## Errors

ACCESS\_ERRORS (*origin*, ATOMIC, MODIFY, APPEND\_LINKS)

If *reverse\_link* is implicit:

ACCESS\_ERRORS (*dest*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

If *reverse\_link* is not implicit:

ACCESS\_ERRORS (*dest*, ATOMIC, MODIFY, APPEND\_LINKS)

If *new\_link* is atomically stabilizing:

ACCESS\_ERRORS (*dest*, ATOMIC, CHANGE, STABILIZE)

If *new\_link* is compositely stabilizing:

ACCESS\_ERRORS (*dest*, COMPOSITE, CHANGE, STABILIZE)

CATEGORY\_IS\_BAD (*origin*, *new\_link*, (COMPOSITION, EXISTENCE, REFERENCE, DESIGNATION))

COMPONENT\_ADDITION\_ERRORS (*dest*, *new\_link*)

COMPONENT\_ADDITION\_ERRORS (*origin*, *reverse\_link*)

DESTINATION\_OBJECT\_TYPE\_IS\_INVALID (*origin*, *new\_link*, *dest*)

LINK\_EXISTS (*origin*, *new\_link*)

If *link* is atomically or compositely stabilizing:

OBJECT\_CANNOT\_BE\_STABILIZED (*dest*)

If *link* is compositely stabilizing:

OBJECT\_CANNOT\_BE\_STABILIZED (component of *dest*)

REVERSE\_KEY\_IS\_BAD (*origin*, *new\_link*, *dest*, *reverse\_key*)

REVERSE\_KEY\_IS\_NOT\_SUPPLIED (*origin*, *new\_link*, *dest*)

REVERSE\_KEY\_IS\_SUPPLIED (*reverse\_key*)

REVERSE\_LINK\_EXISTS (*origin*, *new\_link*, *dest*, *reverse\_key*)

UPPER\_BOUND\_WOULD\_BE\_VIOLATED (*dest*, *reverse\_link*)

UPPER\_BOUND\_WOULD\_BE\_VIOLATED (*origin*, *new\_link*)

USAGE\_MODE\_ON\_LINK\_TYPE\_WOULD\_BE\_VIOLATED (*origin*, *new\_link*, CREATE\_MODE)

## 9.2.2 LINK\_DELETE

```
LINK_DELETE (
    origin : Object_designator,
    link   : Link_designator
)
```

LINK\_DELETE deletes the link specified by *origin* and *link*.

Let *dest* be the destination of *link*, and *reverse\_link* be the reverse link of *link* (if any). LINK\_DELETE deletes *link* from the links of *origin* and deletes *reverse\_link* (if any) from the links of *dest*.

*dest* is deleted from the object base if *link* is the last composition or existence link to *dest*.

*origin* is deleted from the object base if *reverse\_link* is the last composition or existence link to *origin*.

For each deleted object the "object\_on\_volume" link from the volume on which the deleted object was residing to the deleted object is also deleted.

If either deleted object is opened by one or more processes (see 12.1), the deletion of its contents is postponed until all processes have closed the contents. An operation using a contents handle to access its contents is not affected by the deletion until the contents handle is closed.

Write locks of the default mode are obtained on the deleted objects (if any) and on the deleted links except the "object\_on\_volume" link. A read lock of the default mode is obtained on *origin*



if the interpretation of *link* implies the evaluation of any '+' or '++' key attribute values (see 23.1.2.7).

### Errors

ACCESS\_ERRORS (*origin*, ATOMIC, MODIFY, WRITE\_LINKS)

If *link* is atomically stabilizing:

ACCESS\_ERRORS (*dest*, ATOMIC, CHANGE, STABILIZE)

If *link* is compositely stabilizing:

ACCESS\_ERRORS (*dest*, COMPOSITE, CHANGE, STABILIZE)

If *reverse\_link* is implicit:

ACCESS\_ERRORS (*dest*, ATOMIC, CHANGE, WRITE\_IMPLICIT)

If *reverse\_link* is not implicit:

ACCESS\_ERRORS (*dest*, ATOMIC, MODIFY, WRITE\_LINKS)

If *link* is the last composition or existence link to *dest*:

ACCESS\_ERRORS (*dest*, ATOMIC, MODIFY, DELETE)

If *reverse\_link* is the last composition or existence link to *origin*:

ACCESS\_ERRORS (*origin*, ATOMIC, MODIFY, DELETE)

For each origin X of an implicit link to a deleted object:

ACCESS\_ERRORS (X, ATOMIC, CHANGE, WRITE\_IMPLICIT)

For each compositely stabilizing link L of a deleted object:

ACCESS\_ERRORS (destination of L, COMPOSITE, CHANGE, STABILIZE)

CATEGORY\_IS\_BAD (*origin*, *link*, (COMPOSITION, EXISTENCE, REFERENCE, DESIGNATION))

If *link* is not a designation link:

DESTINATION\_OBJECT\_TYPE\_IS\_INVALID (*origin*, *link*, *dest*)

LOWER\_BOUND\_WOULD\_BE\_VIOLATED (*origin*, *link*)

LOWER\_BOUND\_WOULD\_BE\_VIOLATED (*dest*, *reverse\_link*)

If *reverse\_link* is the last existence or composition link to *origin*:

OBJECT\_HAS\_LINKS\_PREVENTING\_DELETION (*origin*)

If *link* is the last existence or composition link to *dest*:

OBJECT\_HAS\_LINKS\_PREVENTING\_DELETION (*dest*)

If *link* is the last composition or existence link to *dest*:

OBJECT\_IS\_IN\_USE\_FOR\_DELETE (*dest*)

If *reverse\_link* is the last composition or existence link to *origin*:

OBJECT\_IS\_IN\_USE\_FOR\_DELETE (*origin*)

USAGE\_MODE\_ON\_LINK\_TYPE\_WOULD\_BE\_VIOLATED (*origin*, *link*, DELETE\_MODE)

### 9.2.3 LINK\_DELETE\_ATTRIBUTE

```
LINK_DELETE_ATTRIBUTE (
    origin      : Object_designator,
    link        : Link_designator,
    attribute   : Attribute_designator
)
```



LINK\_DELETE\_ATTRIBUTE deletes the non-key attribute *attribute* of the link *link* of the object *origin*, if the "attribute\_type" object representing the attribute type of *attribute* is no longer in the object base.

A write lock of the default mode is obtained on *link*. A read lock of the default mode is obtained on *origin* if the interpretation of *link* implies the evaluation of any '+' or '++' key attribute values (see 23.1.2.7).

#### Errors

ACCESS\_ERRORS (*origin*, ATOMIC, MODIFY, WRITE\_LINKS)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_CONFIGURATION)

REFERENCED\_OBJECT\_IS\_NOT\_MUTABLE (key of *link*)

NOTE - It is the responsibility of the user to ensure that the attribute type is no longer in the object base.

### 9.2.4 LINK\_GET\_ATTRIBUTE

```
LINK_GET_ATTRIBUTE (
    origin      : Object_designator,
    link        : Link_designator,
    attribute    : Attribute_designator
)
    result      : Attribute_value
```

LINK\_GET\_ATTRIBUTE returns the value of the non-key attribute *attribute* of the link *link* of the object *origin*.

A read lock of the default mode is obtained on *link*. A read lock of the default mode is obtained on *origin* if the interpretation of *link* implies the evaluation of any '+' or '++' key attribute values (see 23.1.2.7).

#### Errors

ACCESS\_ERRORS (*origin*, ATOMIC, READ, READ\_LINKS)

USAGE\_MODE\_ON\_ATTRIBUTE\_TYPE\_WOULD\_BE\_VIOLATED (*origin*, *link*, *attribute*, READ\_MODE)

### 9.2.5 LINK\_GET\_DESTINATION\_VOLUME

```
LINK_GET_DESTINATION_VOLUME (
    origin      : Object_designator,
    link        : Link_designator
)
    destination : Volume_info
```

LINK\_GET\_DESTINATION\_VOLUME returns the volume identifier *volume\_identifier* and the volume accessibility *mounted* of the volume on which the destination *dest* of the link *link* of the object *origin* resides. The returned value of *mounted* is as follows:

- ACCESSIBLE if the volume on which *dest* resides is mounted and is accessible in the network partition that contains the calling procedure's workstation. In this case, *volume\_identifier* is the volume identifier of the volume on which *dest* resides.

- INACCESSIBLE if the PCTE implementation is able to determine on which volume the object resides, and that volume is not accessible (either because the volume is not mounted or because the volume is mounted in a network partition which does not contain the calling process's workstation). In this case, *volume\_identifier* is the volume identifier of the volume on which the object resides.
- UNKNOWN if the PCTE implementation is unable to determine on which volume the object resides. In this case, *volume\_identifier* is the volume identifier of a volume which is not currently accessible.

The situations in which UNKNOWN is returned rather than INACCESSIBLE, and vice versa, are implementation-defined, as is the general meaning of the volume identifier returned for UNKNOWN. In any particular situation, the choice is implementation-dependent.

Read locks of the default mode are obtained on *link*, and on *origin* if the interpretation of *link* implies the evaluation of any '+' or '++' key attribute values (see 23.1.2.7).

### Errors

ACCESS\_ERRORS (*origin*, ATOMIC, READ, READ\_LINKS)

LINK\_DESTINATION\_DOES\_NOT\_EXIST (*link*)

OBJECT\_IS\_ARCHIVED (destination of *link*)

USAGE\_MODE\_ON\_LINK\_TYPE\_WOULD\_BE\_VIOLATED (*origin*, *link*, NAVIGATE\_MODE)

NOTE - Some implementations may be able to guarantee that only ACCESSIBLE or INACCESSIBLE is returned. For implementations that return UNKNOWN, the volume identifier returned should be that of the volume which should be made accessible prior to repeating the call. The destination object may reside on this volume, or it may contain implementation-dependent details of the volume on which the object resides, or of a further volume to be made accessible. Although the operation may require access to other volumes, no error condition is raised as a result.

## 9.2.6 LINK\_GET\_KEY

```
LINK_GET_KEY (
    origin : Object_designator,
    link   : Link_designator
)
    key   : [ Actual_key ]
```

LINK\_GET\_KEY returns in *key* the complete sequence of key attribute values (if any) of the link *link* of the object *origin*.

Read locks of default mode are obtained on *link*, and on *origin* if the interpretation of *link* implies the evaluation of any '+' or '++' key attribute values (see 23.1.2.7).

### Errors

ACCESS\_ERRORS (*origin*, ATOMIC, READ, READ\_LINKS)

USAGE\_MODE\_ON\_LINK\_TYPE\_WOULD\_BE\_VIOLATED (*origin*, *link*, NAVIGATE\_MODE)

### 9.2.7 LINK\_GET\_REVERSE

```

LINK_GET_REVERSE (
    origin      : Object_designator,
    link        : Link_designator
)
    reverse_link : [ Link_designator ],
    dest         : Object_designator

```

LINK\_GET\_REVERSE returns in *reverse\_link* the reverse link (if there is one) of the link *link* of the object *origin*, and in *dest* the destination of *link*. If *link* has no reverse link, no value is returned in *reverse\_link*.

Read locks of the default mode are obtained on *link*, and on *origin* if the interpretation of *link* implies the evaluation of any '+' or '++' key attribute values (see 23.1.2.7).

#### Errors

ACCESS\_ERRORS (*origin*, ATOMIC, READ, READ\_LINKS)

ACCESS\_ERRORS (destination of *link*, ATOMIC, READ, READ\_LINKS)

LINK\_DESTINATION\_DOES\_NOT\_EXIST (*link*)

LINK\_NAME\_IS\_TOO\_LONG\_IN\_CURRENT\_WORKING\_SCHEMA

REFERENCE\_CANNOT\_BE\_ALLOCATED

USAGE\_MODE\_ON\_LINK\_TYPE\_WOULD\_BE\_VIOLATED (*origin*, *link*, NAVIGATE\_MODE)

### 9.2.8 LINK\_GET\_SEVERAL\_ATTRIBUTES

```

LINK_GET_SEVERAL_ATTRIBUTES (
    origin      : Object_designator,
    link        : Link_designator,
    attributes   : Attribute_selection
)
    values      : Attribute_assignments

```

LINK\_GET\_SEVERAL\_ATTRIBUTES returns in *values* a set of attribute assignments of the link *link* of the object *origin*.

The returned set of attributes is determined by *attributes*:

- a set of attribute designators: the set of non-key attributes, as for LINK\_GET\_ATTRIBUTE (*origin*, *link*, A) for each attribute A of *attributes*;
- VISIBLE\_ATTRIBUTE\_TYPES: all non-key attributes of *link* visible in the working schema of the calling process and with usage mode including READ\_MODE.

Read locks of the default mode are obtained on *link*, and on *origin* if the interpretation of *link* implies the evaluation of any '+' or '++' key attribute values (see 23.1.2.7).

#### Errors

ACCESS\_ERRORS (*origin*, ATOMIC, READ, READ\_LINKS)

If *attributes* is not `VISIBLE_ATTRIBUTE_TYPES`:

`USAGE_MODE_ON_ATTRIBUTE_TYPE_WOULD_BE_VIOLATED` (*origin*,  
*link*, each element of *attributes*, `READ_MODE`)

### 9.2.9 LINK\_REPLACE

```
LINK_REPLACE (
    origin      : Object_designator,
    link        : Link_designator,
    new_origin  : Object_designator,
    new_link    : Link_designator,
    new_reverse_key : [ Actual_key ]
)
```

`LINK_REPLACE` replaces the composition or existence link *link* of the object *origin* by a new composition or existence link as specified by *new\_link* from *new\_origin*, with the same destination *dest*.

The new link from *new\_origin* to *dest* is created and *link* is deleted in the same way as the following sequence of operations, ignoring any temporary violation of link bounds or composition exclusivity on *dest*:

```
LINK_CREATE (new_origin, new_link, dest, new_reverse_key);
```

```
LINK_DELETE (origin, link)
```

Write locks of the default mode are obtained on the links to be deleted and on the new links. A read lock of the default mode is obtained on *new\_origin* if the interpretation of *link* or *new\_link* implies the evaluation of any '+' or '++' key attribute values (see 23.1.2.7). A read lock of the default mode is obtained on *dest* if the interpretation of *new\_reverse\_key* implies the evaluation of any '+' or '++' key attribute values.

A write lock of the default mode is obtained on *dest* and each of its components if the OWNER discretionary access right is granted or denied for one or more groups to *new\_origin*, and different OWNER discretionary access rights exist for one or more of those same groups to *dest*.

#### Errors

```
ACCESS_ERRORS (new_origin, ATOMIC, MODIFY, APPEND_LINKS)
```

```
ACCESS_ERRORS (origin, ATOMIC, MODIFY, WRITE_LINKS)
```

If the reverse link of *new\_link* is implicit:

```
ACCESS_ERRORS (dest, ATOMIC, CHANGE, APPEND_IMPLICIT)
```

If the reverse link of *new\_link* is not implicit:

```
ACCESS_ERRORS (dest, ATOMIC, MODIFY, APPEND_LINKS)
```

If the reverse link of *link* is implicit:

```
ACCESS_ERRORS (dest, ATOMIC, CHANGE, WRITE_IMPLICIT)
```

If the reverse link of *link* is not implicit:

```
ACCESS_ERRORS (dest, ATOMIC, MODIFY, WRITE_LINKS)
```

If *new\_link* is atomically stabilizing and *link* is not, or vice versa:

```
ACCESS_ERRORS (dest, ATOMIC, CHANGE, STABILIZE)
```

If *new\_link* is compositely stabilizing and *link* is not, or vice versa:

```
ACCESS_ERRORS (dest, COMPOSITE, CHANGE, STABILIZE)
```

CATEGORY\_IS\_BAD (new\_origin, new\_link, (COMPOSITION, EXISTENCE))

CATEGORY\_IS\_BAD (origin, link, (COMPOSITION, EXISTENCE))

COMPONENT\_ADDITION\_ERRORS (dest, new\_link)

DESTINATION\_OBJECT\_TYPE\_IS\_INVALID (new\_origin, new\_link, dest.)

DESTINATION\_OBJECT\_TYPE\_IS\_INVALID (origin, link, dest)

If new\_link is of category COMPOSITION, and new\_origin has OWNER granted or denied:

LINK\_EXISTS (new\_origin, new\_link)

LINK\_EXISTS (dest, new\_reverse\_key)

LOWER\_BOUND\_WOULD\_BE\_VIOLATED (origin, link)

OBJECT\_CANNOT\_BE\_STABILIZED (dest)

REVERSE\_KEY\_IS\_BAD (new\_origin, new\_link, new\_reverse\_key)

REVERSE\_KEY\_IS\_NOT\_SUPPLIED (new\_origin, new\_link, dest)

UPPER\_BOUND\_WOULD\_BE\_VIOLATED (dest, reverse link of new\_link)

UPPER\_BOUND\_WOULD\_BE\_VIOLATED (new\_origin, new\_link)

USAGE\_MODE\_ON\_LINK\_TYPE\_WOULD\_BE\_VIOLATED (origin, link, DELETE\_MODE)

USAGE\_MODE\_ON\_LINK\_TYPE\_WOULD\_BE\_VIOLATED (new\_origin, new\_link, CREATE\_MODE)

### 9.2.10 LINK\_RESET\_ATTRIBUTE

```
LINK_RESET_ATTRIBUTE (
    origin      : Object_designator,
    link        : Link_designator,
    attribute    : Attribute_designator
)
```

LINK\_RESET\_ATTRIBUTE resets the non-key attribute *attribute* of the link *link* of the object *origin* to its initial value.

A write lock of the default mode is obtained on *link*. A read lock of the default mode is obtained on *origin* if the interpretation of *link* implies the evaluation of any '+' or '++' key attribute values (see 23.1.2.7).

#### Errors

ACCESS\_ERRORS (*object*, ATOMIC, MODIFY, WRITE\_ATTRIBUTES)

KEY\_UPDATE\_IS\_FORBIDDEN (*attribute*)

REFERENCED\_OBJECT\_IS\_NOT\_MUTABLE (key of *link*)

USAGE\_MODE\_ON\_ATTRIBUTE\_TYPE\_WOULD\_BE\_VIOLATED (*origin*, *link*, *attribute*, WRITE\_MODE)

### 9.2.11 LINK\_SET\_ATTRIBUTE

```
LINK_SET_ATTRIBUTE (
    origin      : Object_designator,
    link        : Link_designator,
    attribute    : Attribute_designator,
    value       : Attribute_value
)
```

LINK\_SET\_ATTRIBUTE assigns the value *value* to the non-key attribute *attribute* of the link *link* of the object *origin*.

A write lock of the default mode is obtained on *link*. A read lock of the default mode is obtained on *origin* if the interpretation of *link* implies the evaluation of any '+' or '++' key attribute values (see 23.1.2.7).

#### Errors

ACCESS\_ERRORS (*origin*, ATOMIC, MODIFY, WRITE\_LINKS)

ENUMERATION\_VALUE\_IS\_OUT\_OF\_RANGE (*value*, values of *attribute*)

KEY\_UPDATE\_IS\_FORBIDDEN (*origin*, *link*, *attribute*)

If *link* is a "referenced\_object" link:

REFERENCED\_OBJECT\_IS\_NOT\_MUTABLE (key of *link*)

USAGE\_MODE\_ON\_ATTRIBUTE\_TYPE\_WOULD\_BE\_VIOLATED (*origin*, *link*, *attribute*, WRITE\_MODE)

VALUE\_LIMIT\_ERRORS (*value*)

The following implementation-dependent error may be raised:

VALUE\_TYPE\_IS\_INVALID (*value*, *origin*, *link*, *attribute*)

### 9.2.12 LINK\_SET\_SEVERAL\_ATTRIBUTES

```
LINK_SET_SEVERAL_ATTRIBUTES (
    origin      : Object_designator,
    link        : Link_designator,
    attributes   : Attribute_assignments
)
```

For each element *A* in the domain of *attributes*, LINK\_SET\_SEVERAL\_ATTRIBUTES sets the value of the attribute *A* of the link *link* of the object *origin* to the value *attributes* (*A*), in the same way as:

LINK\_SET\_ATTRIBUTE (*origin*, *link*, *A*, *attributes* (*A*))

A write lock of the default mode is obtained on *link*. A read lock of the default mode is obtained on *origin* if the interpretation of *link* implies the evaluation of any '+' or '++' key attribute values (see 23.1.2.7).

#### Errors

ACCESS\_ERRORS (*origin*, ATOMIC, MODIFY, WRITE\_LINKS)

For each element *A* in the domain of *attributes*:

ENUMERATION\_VALUE\_IS\_OUT\_OF\_RANGE (*attribute*(*A*), values of *A*)

KEY\_UPDATE\_IS\_FORBIDDEN (*origin*, *link*, *A*)

USAGE\_MODE\_ON\_ATTRIBUTE\_TYPE\_WOULD\_BE\_VIOLATED (*origin*,  
*link*, *A*, WRITE\_MODE)

VALUE\_LIMIT\_ERRORS (*attributes* (*A*))

The following implementation-dependent error may be raised for each element *A* in the domain of *attributes*:

VALUE\_TYPE\_IS\_INVALID (*attributes* (*A*), *origin*, *link*, *A*)

### 9.3 Object operations

#### 9.3.1 OBJECT\_CHECK\_TYPE

```
OBJECT_CHECK_TYPE(
    object      : Object_designator,
    type2       : Object_type_nominator
)
    relation    : Type_ancestry
```

OBJECT\_CHECK\_TYPE compares the object type *type1* of the object *object* against the object type *type2*, and returns in *relation* a value defined as follows:

- EQUAL\_TYPE if *type1* is the same as *type2*;
- ANCESTOR\_TYPE if *type1* is an ancestor of *type2*;
- DESCENDANT\_TYPE if *type1* is a descendant type of *type2* in the working schema of the calling process;
- UNRELATED\_TYPE in all other cases.

The visibility of the type of *object* does not affect the result of the operation.

A read lock of the default mode is obtained on *object*.

#### Errors

ACCESS\_ERRORS (*object*, ATOMIC, READ)

OBJECT\_TYPE\_IS\_UNKNOWN (*type2*)

VOLUME\_IS\_INACCESSIBLE (*object*, ATOMIC)

#### 9.3.2 OBJECT\_CONVERT

```
OBJECT_CONVERT(
    object : Object_designator,
    type   : Object_type_nominator
)
```

OBJECT\_CONVERT changes the object type of the object *object* to the descendant object type *type*.

The operation has no effect if the current type of *object* is already *type*.

A write lock of the default mode is obtained on *object*.

#### Errors

ACCESS\_ERRORS (*object*, ATOMIC, CHANGE, CONTROL\_OBJECT)

OBJECT\_IS\_NOT\_CONVERTIBLE (*object*)

OBJECT\_IS\_STABLE (*object*)

OBJECT\_TYPE\_IS\_INVALID (*type*)

OBJECT\_TYPE\_IS\_UNKNOWN (*type*)

TYPE\_IS\_NOT\_DESCENDANT (object type of *object*, *type*)

If *object* is not of type *type*:

USAGE\_MODE\_ON\_OBJECT\_TYPE\_WOULD\_BE\_VIOLATED (current type of *object*, *type*)

### 9.3.3 OBJECT\_COPY

```
OBJECT_COPY (
    object           : Object_designator,
    new_origin      : Object_designator,
    new_link        : Link_designator,
    reverse_key     : [ Actual_key ],
    on_same_volume_as : [ Object_designator ],
    access_mask     : Atomic_access_rights
)
    new_object      : Object_designator
```

OBJECT\_COPY creates an object *new\_object* as a copy of *object*. More precisely:

- If *object* is a file, an accounting log, or an audit file, then the contents of *new\_object* is the same as the contents of *object*; if *object* is a pipe then the contents of *new\_object* is empty.
- For each duplicable attribute X of *object*, there is an attribute of *new\_object* which is a copy of X.
- For each duplicable direct component X of *object*, there is a direct component of *new\_object* which is a copy of X.
- For each duplicable internal link A of *object* whose destination is a duplicable component, there is an internal link B of *new\_object* such that the destination of B is the copy of the destination of A and all other properties of B are the same as for A.
- For each duplicable external link A of *object*, there is a corresponding external link of *new\_object* which is a copy of A.
- For each non-duplicable attribute of *object*, there exists a corresponding attribute of *new\_object* whose value is either set to the initial value of the attribute type or, for the following predefined attributes: the exact identifier, volume identifier, replicated state, contents type, last access time, last modification time, last change time, last composite access time, last composite modification time, last composite change time, number of incoming links, number of incoming composition links, number of incoming existence links, number of incoming reference links, number of incoming stabilizing links, number of outgoing composition links, number of outgoing existence links, atomic ACL, composite ACL, confidentiality label, and integrity label, is set to a value corresponding to the newly created object. In particular, since the attribute "replicated\_state" is not copied, the copy of a replicated object is not replicated.
- For each non-duplicable non-key attribute of a copied link, there exists a corresponding non-key attribute of the copied link whose value is set to the initial value of the attribute type.



A *copy* of an attribute *A* is an attribute which has the same attribute type, attribute value, and attribute properties as *A*.

A *copy* of a link *A* is a link which has the same link type, key attributes, duplicable non-key attributes, link properties, and destination as *A*.

A *copy* of a component *X* of an object *A* is a component *Y* of a copy *B* of *A* such that *Y* is a copy of *X* as an object, and the composition link from *B* to *Y* is a copy of the composition link from *X* to *A*.

OBJECT\_COPY creates a link *link* of the object *new\_origin*, as specified by *new\_link*, and with *new\_object* as destination, and its reverse link *reverse\_link* with origin *new\_object* and key derived from *reverse\_key* as described in 23.1.2.7.

*access\_mask* is used in conjunction with the default atomic ACL and default object owner of the calling process to specify the atomic ACL and the composite ACL of *new\_object* and its components. See 19.1.4 for more details.

If *new\_link* is a composition link, then any security group that has OWNER granted or denied to *new\_origin* has OWNER granted or denied respectively to *dest*; similarly if *reverse\_link* is a composition link, then any security group that has OWNER granted or denied to *dest* has OWNER granted or denied respectively to *new\_origin*.

If *new\_link* is a composition link, then any security group that has OWNER granted or denied to *origin* has OWNER granted or denied respectively to *dest*; similarly if *reverse\_link* is a composition link, then any security group that has OWNER granted or denied to *dest* has OWNER granted or denied respectively to *origin*.

*new\_object* has the same integrity and confidentiality label as *object* and each component of *new\_object* has the same integrity and confidentiality labels as the corresponding component of *object*.

If *on\_same\_volume\_as* is supplied, *new\_object* resides on the same volume as the object *on\_same\_volume\_as*. Otherwise, *new\_object* resides on the same volume as *object* and each component of *new\_object* resides on the same volume as its corresponding component in *object*.

An "object\_on\_volume" link is created from the volume on which *new\_object* resides to *new\_object*, and similarly for all its components. Each created link is keyed by the exact identifier of its destination object.

Read locks of the default mode are obtained on *object* and on all its components; write locks of the default mode are obtained on the new objects and links except the new "object\_on\_volume" links. A read lock of the default mode is obtained on *new\_object* if the interpretation of the link *new\_link* implies the evaluation of any '+' or '++' key attribute values (see 23.1.2.7).

If *object* is an accounting log, its contents is preserved.

## Errors

ACCESS\_ERRORS (*new\_origin*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*object*, COMPOSITE, READ, READ\_LINKS)

ACCESS\_ERRORS (*object*, COMPOSITE, READ, READ\_ATTRIBUTES)

ACCESS\_ERRORS (*object*, COMPOSITE, READ, READ\_CONTENTS)

ACCESS\_ERRORS (*on\_same\_volume\_as*, ATOMIC, SYSTEM\_ACCESS)

If *new\_link* is compositely stabilizing:

ACCESS\_ERRORS (*new\_object*, COMPOSITE, CHANGE, STABILIZE)

If *reverse\_link* is compositely stabilizing:

ACCESS\_ERRORS (*new\_origin*, COMPOSITE, CHANGE, STABILIZE)

For each destination X of a duplicable external link L of a duplicated component:

ACCESS\_ERRORS (X, ATOMIC, CHANGE, APPEND\_IMPLICIT)

If L is atomically stabilizing:

ACCESS\_ERRORS (X, ATOMIC, CHANGE, STABILIZE)

If L is compositely stabilizing:

ACCESS\_ERRORS (X, COMPOSITE, CHANGE, STABILIZE)

CATEGORY\_IS\_BAD (*new\_origin*, *new\_link*, (COMPOSITION, EXISTENCE))

If *object* is a component of itself, and *new\_link* is a composition link:

COMPONENT\_ADDITION\_ERRORS (*new\_object*, *new\_link*)

CONTROL\_WOULD\_NOT\_BE\_GRANTED (*new\_object*)

DESTINATION\_OBJECT\_TYPE\_IS\_INVALID (*new\_origin*, *new\_link*, *new\_object*)

EXTERNAL\_LINK\_IS\_BAD (*object*, COMPOSITE)

EXTERNAL\_LINK\_IS\_NOT\_DUPLICABLE (*object*)

LABEL\_IS\_OUTSIDE\_RANGE (*new\_object*, volume on which *on\_same\_volume\_as* resides)

LINK\_EXISTS (*new\_origin*, *new\_link*)

OBJECT\_OWNER\_VALUE\_WOULD\_BE\_INCONSISTENT\_WITH\_ATOMIC\_ACL  
(*new\_object*)

If *object* is not a component of itself, *new\_link* is a composition link, and *new\_origin* has  
OWNER granted or denied:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION (*new\_object*)

REFERENCE\_CANNOT\_BE\_ALLOCATED

REVERSE\_KEY\_IS\_BAD (*new\_origin*, *new\_link*, *new\_object*, *reverse\_key*)

REVERSE\_KEY\_IS\_NOT\_SUPPLIED (*new\_origin*, *new\_link*, *new\_object*)

REVERSE\_KEY\_IS\_SUPPLIED (*reverse\_key*)

REVERSE\_LINK\_EXISTS (*new\_origin*, *new\_link*, *new\_object*, *reverse\_key*)

TYPE\_OF\_OBJECT\_IS\_INVALID (*object*, COMPOSITE)

USAGE\_MODE\_ON\_LINK\_TYPE\_WOULD\_BE\_VIOLATED (*new\_origin*, *new\_link*,  
CREATE)

USAGE\_MODE\_ON\_OBJECT\_TYPE\_WOULD\_BE\_VIOLATED ("object", type of *object*)

VOLUME\_IS\_FULL (volume on which *on\_same\_volume\_as* resides)

NOTE - Key values of reverse links (which must be implicit of cardinality many) are system-generated, because they must be different from key values of other links of the same types from the same origins.

### 9.3.4 OBJECT\_CREATE

```

OBJECT_CREATE (
    type           : Object_type_nominator,
    new_origin     : Object_designator,
    new_link       : Link_designator,
    reverse_key    : [ Actual_key ],
    on_same_volume_as : [ Object_designator ],
    access_mask    : Atomic_access_rights
)
    new_object     : Object_designator

```

OBJECT\_CREATE creates an object *new\_object* as follows:

- the object type of *new\_object* is *type*;
- the contents of *new\_object* is empty;
- the value of each attribute of *new\_object* is the initial value of its attribute type, except for some predefined attributes, set as defined below.

A composition or existence link, as specified by *new\_link*, is created from *new\_origin* to *new\_object*, together with its reverse link *reverse\_link*, with key *reverse\_key*.

*access\_mask* is used in conjunction with the default atomic ACL and default object owner of the calling process to specify the atomic ACL and the composite ACL of *new\_object*. See 19.1.4 for more details.

If *new\_link* is a composition link, then any security group that has OWNER granted or denied to *new\_origin* has OWNER granted or denied respectively to *new\_object*; similarly if *reverse\_link* is a composition link, then any security group that has OWNER granted or denied to *new\_object* has OWNER granted or denied respectively to *new\_origin*.

The confidentiality label of *new\_object* is set to the current confidentiality context of the calling process and the integrity label of *new\_object* is set to the current integrity context of the calling process.

If *on\_same\_volume\_as* is supplied, *new\_object* resides on the same volume as the object *on\_same\_volume\_as*. Otherwise, *new\_object* resides on the same volume as *new\_origin*.

An "object\_on\_volume" link is created from the volume on which *new\_object* resides to *new\_object*. Each created link is keyed by the exact identifier of its destination object.

Write locks of the default mode are obtained on *new\_object* and on *new\_link*. A read lock of the default mode is obtained on *new\_origin* if the interpretation of *new\_link* implies the evaluation of any '+' or '++' key attribute values (see 23.1.2.7).

#### Errors

ACCESS\_ERRORS (*new\_origin*, ATOMIC, MODIFY, APPEND\_LINKS)

If *new\_link* is atomically stabilizing:

ACCESS\_ERRORS (*new\_object*, ATOMIC, CHANGE, STABILIZE)

If *new\_link* is compositely stabilizing:

ACCESS\_ERRORS (*new\_object*, COMPOSITE, CHANGE, STABILIZE)

If *reverse\_link* is compositely stabilizing:

ACCESS\_ERRORS (*new\_origin*, COMPOSITE, CHANGE, STABILIZE)

ACCESS\_ERRORS (*on\_same\_volume\_as*, ATOMIC, SYSTEM\_ACCESS)  
 CATEGORY\_IS\_BAD (*new\_origin*, *new\_link*, (COMPOSITION, EXISTENCE)  
 CONTROL\_WOULD\_NOT\_BE\_GRANTED (*new\_object*)  
 DESTINATION\_OBJECT\_TYPE\_IS\_INVALID (*new\_origin*, *new\_link*, *new\_object*)  
 LABEL\_IS\_OUTSIDE\_RANGE (*new\_object*, volume on which *on\_same\_volume\_as* resides)  
 LINK\_EXISTS (*new\_origin*, *new\_link*)  
 OBJECT\_OWNER\_VALUE\_WOULD\_BE\_INCONSISTENT\_WITH\_ATOMIC\_ACL (*new\_object*)  
 OBJECT\_TYPE\_IS\_INVALID (*type*)  
 OBJECT\_TYPE\_IS\_UNKNOWN (*type*)  
 If *new\_link* is a composition link and *new\_origin* has OWNER granted or denied:  
     OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION (*new\_object*)  
 If *reverse\_link* is a composition link and *new\_object* has OWNER granted or denied:  
     OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION (*new\_object*)  
 REFERENCE\_CANNOT\_BE\_ALLOCATED  
 REVERSE\_KEY\_IS\_NOT\_SUPPLIED (*new\_origin*, *new\_link*, *new\_object*)  
 REVERSE\_KEY\_IS\_BAD (*new\_origin*, *new\_link*, *new\_object*, *reverse\_key*)  
 REVERSE\_KEY\_IS\_SUPPLIED (*reverse\_key*)  
 UPPER\_BOUND\_WOULD\_BE\_VIOLATED (*new\_origin*, *new\_link*)  
 USAGE\_MODE\_ON\_LINK\_TYPE\_WOULD\_BE\_VIOLATED (*new\_origin*, *new\_link*,  
 CREATE\_MODE)  
 USAGE\_MODE\_ON\_OBJECT\_TYPE\_WOULD\_BE\_VIOLATED ("object", *type*)

### 9.3.5 OBJECT\_DELETE

```

OBJECT_DELETE (
    origin : Object designator,
    link   : Link designator
)
  
```

OBJECT\_DELETE deletes the composition or existence link *link* of the object *origin*, its reverse link *reverse\_link*, and possibly its destination *dest*. More precisely:

- the link *link* of the object *origin* is deleted; the reverse link *reverse\_link* of *link* is deleted;
- if *link* is the last existence or composition link to *dest*, then *dest* is deleted.

To *delete* an object X entails the deletion of all components of X except components Y for which there is an incoming external link with the existence property to Y or to an enclosing object of Y, and the deletion of all links from and to those deleted objects (except designation links to them).

Non-implicit links of components which are not deleted are not affected.

If *origin* or any of its components is opened by one or more processes (see 12.1), the deletion of its contents is postponed until all processes have closed the contents: i.e. the object is no longer accessible but an operation using a contents handle to access its contents is not affected by the deletion until the contents handle is closed.

For each deleted object, if any, the "object\_on\_volume" link from the volume on which the deleted object resided to the deleted object is also deleted.

Write locks of the default kind are obtained on the deleted objects, if any, and on the deleted links except the deleted "object\_on\_volume" links; and a read lock of the default mode is obtained on *origin* if the interpretation of *link* implies the evaluation of any '+' or '++' key attribute values (see 23.1.2.7).

### Errors

If the conditions hold for the deletion of *dest*:

ACCESS\_ERRORS (*dest* and its deleted components, ATOMIC, MODIFY, DELETE)

ACCESS\_ERRORS (*origin*, ATOMIC, MODIFY, WRITE\_LINKS)

If *reverse\_link* is not implicit:

ACCESS\_ERRORS (*dest*, ATOMIC, MODIFY, WRITE\_LINKS)

If *reverse\_link* is implicit:

ACCESS\_ERRORS (*dest*, ATOMIC, CHANGE, WRITE\_IMPLICIT)

For each origin X of an implicit incoming external link to a deleted object:

ACCESS\_ERRORS (X, ATOMIC, CHANGE, WRITE\_IMPLICIT)

For each atomically stabilizing external link L of a deleted object:

ACCESS\_ERRORS (destination of L, ATOMIC, CHANGE, STABILIZE)

For each compositely stabilizing external link L of a deleted object:

ACCESS\_ERRORS (destination of L, COMPOSITE, CHANGE, STABILIZE)

CATEGORY\_IS\_BAD (*origin*, *link*, (EXISTENCE, COMPOSITION))

DESTINATION\_OBJECT\_TYPE\_IS\_INVALID (*origin*, *link*, *dest*)

LOWER\_BOUND\_WOULD\_BE\_VIOLATED (*origin*, *link*)

OBJECT\_HAS\_EXTERNAL\_LINKS\_PREVENTING\_DELETION (*dest*)

OBJECT\_HAS\_INTERNAL\_LINKS\_PREVENTING\_DELETION (*dest*)

OBJECT\_IS\_IN\_USE\_FOR\_DELETE (*dest*)

USAGE\_MODE\_ON\_LINK\_TYPE\_WOULD\_BE\_VIOLATED (*origin*, *link*, DELETE\_MODE)

NOTE - OBJECT\_DELETE works exactly like LINK\_DELETE if *dest* has no direct components.

### 9.3.6 OBJECT\_DELETE\_ATTRIBUTE

```
OBJECT_DELETE_ATTRIBUTE (
    object      : Object_designator,
    attribute    : Attribute_designator
)
```

OBJECT\_DELETE\_ATTRIBUTE removes the attribute *attribute* from the attributes of the object *object*, if the "attribute\_type" object representing the attribute type of *attribute* is no longer in the object base.

A write lock of the default mode is obtained on *object*.

### Errors

ACCESS\_ERRORS (*object*, ATOMIC, MODIFY, WRITE\_ATTRIBUTES)

## PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_CONFIGURATION)

NOTE - It is the responsibility of the user to ensure that the attribute type is no longer in the object base.

**9.3.7 OBJECT\_GET\_ATTRIBUTE**

```
OBJECT_GET_ATTRIBUTE (
    object      : Object_designator,
    attribute    : Attribute_designator
)
    value      : Attribute_value
```

OBJECT\_GET\_ATTRIBUTE returns the value *value* of the attribute *attribute* of the object *object*.

A read lock of the default mode is obtained on *object*. If *attribute* is a predefined composite time, a read lock of the default mode is also obtained on all components of *object*.

**Errors**

ACCESS\_ERRORS (*object*, ATOMIC, READ, READ\_ATTRIBUTES)

If *attribute* is the last composite access time, last composite modification time, or last composite change time:

ACCESS\_ERRORS (*object*, COMPOSITE, READ, READ\_ATTRIBUTES)

USAGE\_MODE\_ON\_ATTRIBUTE\_TYPE\_WOULD\_BE\_VIOLATED (*object*, *attribute*, READ\_MODE)

**9.3.8 OBJECT\_GET\_PREFERENCE**

```
OBJECT_GET_PREFERENCE (
    object : Object_designator
)
    key    : [ Text ],
    type   : [ Link_type_nominator ]
```

OBJECT\_GET\_PREFERENCE returns the preferred link key *key* and preferred link type *type*, if any, of the object *object*.

A read lock of the default mode is obtained on *object*.

**Errors**

ACCESS\_ERRORS (*object*, ATOMIC, READ, READ\_ATTRIBUTES)

**9.3.9 OBJECT\_GET\_SEVERAL\_ATTRIBUTES**

```
OBJECT_GET_SEVERAL_ATTRIBUTES (
    object      : Object_designator,
    attributes    : Attribute_selection
)
    values      : Attribute_assignments
```

OBJECT\_GET\_SEVERAL\_ATTRIBUTES returns a set of attribute assignments *values* of the object *object*.

The returned set of attributes is determined by *attributes*:

- a set of attribute designators: the set of attributes, as for OBJECT\_GET\_ATTRIBUTE (*object*, *A*) for each attribute *A* of *attributes*;
- VISIBLE\_ATTRIBUTE\_TYPES: all attributes of *object* visible in the working schema of the calling process and with usage mode including READ\_MODE.

A read lock of the default mode is obtained on *object*. If any of the attributes is a predefined composite time, a read lock of the default mode is also obtained on all components of *object*.

#### Errors

ACCESS\_ERRORS (*object*, ATOMIC, READ, READ\_ATTRIBUTES)

If *attributes* contains one or more of last composite access time, last composite modification time, or last composite change time, or if *attributes* is VISIBLE\_ATTRIBUTE\_TYPES and one or more of those attributes are visible in the working schema of the calling process with usage mode READ\_MODE:

ACCESS\_ERRORS (*object*, COMPOSITE, READ, READ\_ATTRIBUTES)

USAGE\_MODE\_ON\_ATTRIBUTE\_TYPE\_WOULD\_BE\_VIOLATED (*object*, an element of *attributes*, READ\_MODE)

### 9.3.10 OBJECT\_GET\_TYPE

```
OBJECT_GET_TYPE (
    object : Object_designator
)
    type : Object_type_nominator
```

OBJECT\_GET\_TYPE returns the type *type* of the object *object*; i.e. the actual type of *object*, whether or not it is visible.

A read lock of the default mode is obtained on *object*.

#### Errors

CONFIDENTIALITY\_WOULD\_BE\_VIOLATED (*object*)

INTEGRITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED (*object*)

OBJECT\_IS\_ARCHIVED (*object*)

VOLUME\_IS\_INACCESSIBLE (*object*, ATOMIC)

### 9.3.11 OBJECT\_IS\_COMPONENT

```
OBJECT_IS_COMPONENT (
    object1 : Object_designator,
    object2 : Object_designator
)
    value : Boolean
```

OBJECT\_IS\_COMPONENT tests if *object1* is a component of *object2*.

If *object1* is a component of *object2*, *value* is **true**, otherwise it is **false**.

Read locks of the default mode are obtained on *object1* and on *object2*, and on accessed components of *object2*.



**Errors**ACCESS\_ERRORS (*object2*, COMPONENTS, READ, READ\_LINKS)**9.3.12 OBJECT\_LIST\_LINKS**

```

OBJECT_LIST_LINKS (
    origin      : Object_designator,
    extent      : Link_scope,
    scope       : Object_scope,
    categories : [ Categories ],
    visibility  : Link_selection
)
    links      : Link_set_descriptors

```

OBJECT\_LIST\_LINKS returns in *links* a set of links of the object *origin* and possibly of its components determined by *extent*, *scope*, *categories*, and *visibility*.

*extent* affects the returned set of links as follows:

- INTERNAL\_LINKS: only internal links are returned.
- EXTERNAL\_LINKS: only external links are returned.
- ALL\_LINKS: both internal and external links are returned.

In the lists of links returned, designation links to deleted objects appear only when *extent* is ALL\_LINKS or EXTERNAL\_LINKS.

*scope* affects the returned set of links as follows:

- ATOMIC: only links of *origin* are returned.
- COMPOSITE: links of *origin* and of all components of *origin* are returned.

*categories* may be omitted if *visibility* is a set of link type nominators, and is ignored in that case if supplied. In other cases only link types with category in the set *categories* are returned.

*visibility* restricts the returned set of links as follows:

- VISIBLE\_LINK\_TYPES: only links with link type which is visible in the calling process's working schema are returned;
- ALL\_LINK\_TYPES: all links are returned;
- a set of link type nominators: only links with link type identified by an element of the set are returned.

A read lock of the default mode is obtained on *origin*, and if *scope* is COMPOSITE read locks of the default mode are obtained on all its components.

**Errors**ACCESS\_ERRORS (*origin*, *scope*, READ, READ\_LINKS)

LINK\_NAME\_IS\_TOO\_LONG\_IN\_CURRENT\_WORKING\_SCHEMA

REFERENCE\_CANNOT\_BE\_ALLOCATED



## NOTES

- 1 When *scope* is ATOMIC or *origin* has no components, the object designator in each returned link designates the object *origin*. In other cases, the designated objects may include the object *origin* and its components.
- 2 Each object designator returned in a link set descriptor of *links* can be used with each link designator of that link set descriptor to retrieve information about the volume on which the object resides, by means of LINK\_GET\_DESTINATION\_VOLUME.
- 3 If *scope* is COMPOSITE, links of *origin* and of all components of *origin* are returned. It is possible that the origins of some of the returned links are not accessible from *origin* through paths of visible links.
- 4 OBJECT\_LIST\_LINKS does not prevent a process from seeing data structures which are inconsistent with the visibility restrictions of its working schema.

**9.3.13 OBJECT\_LIST\_VOLUMES**

```
OBJECT_LIST_VOLUMES(
    object      : Object_designator
)
    volumes    : Volume_infos
```

OBJECT\_LIST\_VOLUMES returns the set of volume identifiers of volumes holding components of *object* (except for any components to which there are composition links from components on unmounted volumes), with an indication of the mounted state of each volume.

A read lock of the default mode is obtained on *object*.

**Errors**

ACCESS\_ERRORS (*object*, COMPOSITE, READ, READ\_LINKS)

OBJECT\_IS\_ARCHIVED (component of *object*)

USAGE\_MODE\_ON\_LINK\_TYPE\_WOULD\_BE\_VIOLATED (component of *object*, direct outgoing composition link of that component, NAVIGATE\_MODE)

NOTE - The note of 9.2.5 applies for each component of *object*.

**9.3.14 OBJECT\_MOVE**

```
OBJECT_MOVE (
    object      : Object_designator,
    on_same_volume_as : Object_designator,
    scope       : Object_scope
)
```

OBJECT\_MOVE moves *object* to the volume *volume* on which *on\_same\_volume\_as* resides.

If *scope* is ATOMIC:

- If *object* already resides on *volume*, it is not affected.
- Otherwise, *object* is moved to *volume*.

If *scope* is COMPOSITE:

- Each of *object* and the components of *object* which already reside on *volume* are not affected.

- All other of *object* and the components of *object* are moved to *volume*, and the space previously occupied by those components is freed.

The effect of *moving* an object *A* to a volume *V* is as follows.

- The attributes and links of *A* are unchanged, except for the predefined attributes "volume\_identifier" which is set to the volume identifier of *V*, and "last\_change\_time" and "last\_composite\_change\_time", which are set to the current system time.
- For *object* (if moved) and each moved component, the "object\_on\_volume" link to it from the volume on which the component was previously residing is deleted, and a new "object\_on\_volume" link is created to it from *volume*. The created link is keyed by the exact identifier of its destination object.

A write lock of the default mode is obtained on each moved object. An implementation may set a write lock of the default mode on each link to a moved object (except the "object\_on\_volume" links) if the link is modified by the operation.

### Errors

If *scope* is ATOMIC:

ACCESS\_ERRORS (*object*, ATOMIC, CHANGE, CONTROL\_OBJECT)

If *scope* is COMPOSITE:

ACCESS\_ERRORS (*object*, COMPOSITE, READ, READ\_LINKS)

ACCESS\_ERRORS (*object* and its moved components, ATOMIC, CHANGE, CONTROL\_OBJECT)

ACCESS\_ERRORS (*on\_same\_volume\_as*, ATOMIC, SYSTEM\_ACCESS)

If *object* or some of its components are moved:

OBJECT\_IS\_IN\_USE\_FOR\_MOVE (*object*)

OBJECT\_IS\_INACCESSIBLY\_ARCHIVED (*object*, *scope*)

OBJECT\_IS\_LOCKED (*object*, *scope*)

OBJECT\_IS\_NOT\_MOVABLE (*object*, *scope*)

OBJECT\_IS\_REPLICATED (*object*, *scope*)

TYPE\_OF\_OBJECT\_IS\_INVALID (*object*, *scope*)

VOLUME\_IS\_FULL (*volume* of *on\_same\_volume\_as*)

The following implementation-dependent errors may be raised for any object *X* with a link to *object*:

OBJECT\_IS\_INACCESSIBLY\_ARCHIVED (*X*)

VOLUME\_IS\_INACCESSIBLE (*volume* on which *X* resides)

VOLUME\_IS\_READ\_ONLY (*volume* on which *X* resides)

### 9.3.15 OBJECT\_RESET\_ATTRIBUTE

```
OBJECT_RESET_ATTRIBUTE (
    object      : Object_designator,
    attribute   : Attribute_designator
)
```

OBJECT\_RESET\_ATTRIBUTE resets the attribute *attribute* of the object *object* to its initial value.

A write lock of the default mode is obtained on *object*.

**Errors**

ACCESS\_ERRORS (*object*, ATOMIC, MODIFY, WRITE\_ATTRIBUTES)

USAGE\_MODE\_ON\_ATTRIBUTE\_TYPE\_WOULD\_BE\_VIOLATED (*object*, *attribute*, WRITE\_MODE)

**9.3.16 OBJECT\_SET\_ATTRIBUTE**

```
OBJECT_SET_ATTRIBUTE (
    object      : Object_designator,
    attribute   : Attribute_designator,
    value       : Attribute_value
)
```

OBJECT\_SET\_ATTRIBUTE assigns the value *value* to the attribute *attribute* of the object *object*.

A write lock of the default mode is obtained on *object*.

**Errors**

ACCESS\_ERRORS (*object*, ATOMIC, MODIFY, WRITE\_ATTRIBUTES)

ENUMERATION\_VALUE\_IS\_OUT\_OF\_RANGE (*value*, values of *attribute*)

USAGE\_MODE\_ON\_ATTRIBUTE\_TYPE\_WOULD\_BE\_VIOLATED (*object*, *attribute*, WRITE\_MODE)

VALUE\_LIMIT\_ERRORS (*value*)

The following implementation-dependent error may be raised:

VALUE\_TYPE\_IS\_INVALID (*value*, *object*, *attribute*)

**9.3.17 OBJECT\_SET\_PREFERENCE**

```
OBJECT_SET_PREFERENCE (
    object : Object_designator,
    type   : [ Link_type_nominator ],
    key    : [ Text ]
)
```

OBJECT\_SET\_PREFERENCE sets the preferred link type of the object *object* to the link type *type* (if supplied), and preferred link key of *object* to *key* (if supplied).

If both *type* and *key* are supplied, the preferred link type of *object* is set to *type* and the preferred link key of *object* is set to *key*.

If *type* is supplied and *key* is not, the preferred link type of *object* is set to *type* and the preferred link key of *object* is unset.

If *type* is not supplied and *key* is supplied, then the preferred link type of *object* must already be set (else the error condition PREFERRED\_LINK\_TYPE\_IS\_UNSET is raised); the preferred link key is set to *key*.

If *key* and *type* are not supplied, the preferred link type and preferred link key of *object* are unset.

A write lock of the default mode is obtained on *object*.

**Errors**

ACCESS\_ERRORS (*object*, ATOMIC, MODIFY, WRITE\_ATTRIBUTES)

CARDINALITY\_IS\_INVALID (*type*)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_KEY\_SIZE)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_KEY\_VALUE)

LINK\_TYPE\_IS\_UNKNOWN (*type*)

PREFERRED\_LINK\_KEY\_IS\_BAD (*key*, *type* or preferred link type of *object* if *type* is not supplied)

If *type* is not supplied and *key* is supplied:

PREFERRED\_LINK\_TYPE\_IS\_UNSET (*object*)

**9.3.18 OBJECT\_SET\_SEVERAL\_ATTRIBUTES**

```
OBJECT_SET_SEVERAL_ATTRIBUTES (
    object      : Object_designator,
    attributes  : Attribute_assignments
)
```

For each element A of the domain of *attributes*, OBJECT\_SET\_SEVERAL\_ATTRIBUTES sets the value of the attribute A of *object* to *attributes* (A) in the same way as:

OBJECT\_SET\_ATTRIBUTE (*object*, A, *attributes* (A)).

A write lock of the default mode is obtained on *object*.

**Errors**

ACCESS\_ERRORS (*object*, ATOMIC, MODIFY, WRITE\_ATTRIBUTES)

For each element A of the domain of *attributes*

ENUMERATION\_VALUE\_IS\_OUT\_OF\_RANGE (*attributes*(A), values of A)

USAGE\_MODE\_ON\_ATTRIBUTE\_TYPE\_WOULD\_BE\_VIOLATED (*object*, A, WRITE\_MODE)

VALUE\_LIMIT\_ERRORS (*attributes* (A))

The following implementation-dependent error may be raised for each element A of the domain of *attributes*:

VALUE\_TYPE\_IS\_INVALID (*attributes* (A), *object*, A)

**9.3.19 OBJECT\_SET\_TIME\_ATTRIBUTES**

```
OBJECT_SET_TIME_ATTRIBUTES(
    object      : Object_designator,
    last_access : [ Time ],
    last_modification : [ Time ],
    scope      : Object_scope
)
```

OBJECT\_SET\_TIME\_ATTRIBUTES sets the time attributes of the object *object* as follows.

If *scope* is ATOMIC:

- the last access time of *object* is set to *last\_access* if supplied, otherwise to the current system time;

- the last modification time of *object* is set to *last\_modification* if supplied, otherwise to the current system time;

If *scope* is COMPOSITE:

- the last composite access time of *object*, and the last access time of each component of *object*, are set to *last\_access* if supplied, otherwise to the current system time;
- the last composite modification time of *object*, and the last modification time of each component of *object*, are set to *last\_modification* if supplied, otherwise to the current system time;

A write lock of the default mode is obtained on *object*.

#### Errors

ACCESS\_ERRORS (*object*, *scope*, MODIFY, WRITE\_ATTRIBUTES)

If *last\_access* or *last\_modification* is supplied:

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_HISTORY)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_TIME\_ATTRIBUTE, MIN\_TIME\_ATTRIBUTE)

### 9.3.20 VOLUME\_LIST\_OBJECTS

```
VOLUME_LIST_OBJECTS (
    volume      : Volume_designator,
    types       : Object_type_nominators
)
    objects     : Object_designators
```

VOLUME\_LIST\_OBJECTS returns in *objects* a set of object designators determined by *types*.

An object designator is returned in *objects* for each object which resides on *volume*, whose type in working schema is an element of *types*.

A read lock of the default mode is obtained on *volume*.

#### Errors

ACCESS\_ERRORS (*volume*, ATOMIC, READ, READ\_LINKS)

REFERENCE\_CANNOT\_BE\_ALLOCATED

## 9.4 Version operations

### 9.4.1 VERSION\_ADD\_PREDECESSOR

```
VERSION_ADD_PREDECESSOR (
    version      : Object_designator,
    new_predecessor : Object_designator
)
```

VERSION\_ADD\_PREDECESSOR adds *new\_predecessor* as a predecessor of *version* in a graph of versions, by creating a "predecessor" link with key the next available natural value from *version* to *new\_predecessor*.

Write locks of the default mode are obtained on the new links.

**Errors**

ACCESS\_ERRORS (*new\_predecessor*, ATOMIC, CHANGE, STABILIZE)

ACCESS\_ERRORS (*version*, ATOMIC, MODIFY, APPEND\_LINKS)

MASTER\_IS\_INACCESSIBLE (some object of the graph of security groups, ATOMIC)

OBJECT\_CANNOT\_BE\_STABILIZED (component of *version*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_HISTORY)

VERSION\_GRAPH\_IS\_INVALID (*version*, *new\_predecessor*)

**9.4.2 VERSION\_IS\_CHANGED**

```

VERSION_IS_CHANGED (
    version      : Object_designator,
    predecessor : Natural
)
    changed      : Boolean
  
```

VERSION\_IS\_CHANGED tests whether *version* has been changed since being created as a new version of its predecessor *predecessor* by comparing the values of the last composite modification time for *version* and *predecessor*. If it has been changed, i.e. the last composite modification times are different, then *changed* is **true**, otherwise it is **false**.

Read locks of the default mode are obtained on *version*, on all components of *version*, and on *predecessor*.

**Errors**

ACCESS\_ERRORS (*predecessor*, COMPONENTS, READ, (READ\_LINKS, READ\_ATTRIBUTES, READ\_CONTENTS))

ACCESS\_ERRORS (*version*, COMPONENTS, READ, (READ\_LINKS, READ\_ATTRIBUTES, READ\_CONTENTS))

ACCESS\_ERRORS (*version*, COMPOSITE, CHANGE)

LINK\_DOES\_NOT\_EXIST (*version*, "predecessor" link with key *predecessor*)

**9.4.3 VERSION\_REMOVE**

```

VERSION_REMOVE (
    version      : Object_designator
)
  
```

VERSION\_REMOVE removes *version* from its graph of versions.

For X as *version* and each component of *version*:

- If X has external successors and predecessors then a "predecessor" link is created from each successor of X to each predecessor of X, and all the "predecessor" links to and from X are deleted.
- If X has only external successors: then all the "predecessor" links to X are deleted.
- If X has only external predecessors then all the "predecessor" links from X are deleted.

Write locks of the default mode are obtained on the deleted links and on the created links.

**Errors**

ACCESS\_ERRORS (*version*, COMPOSITE, MODIFY, (WRITE\_LINKS, WRITE\_IMPLICIT))

For each component X of *version* which has more than one successor:

ACCESS\_ERRORS (predecessor of X, COMPOSITE, CHANGE, STABILIZE)

ACCESS\_ERRORS (predecessor of *version*, ATOMIC, CHANGE, (WRITE\_IMPLICIT, APPEND\_IMPLICIT))

ACCESS\_ERRORS (successor of *version*, ATOMIC, CHANGE, (WRITE\_LINKS, APPEND\_LINKS))

ACCESS\_ERRORS (*version* and its deleted components, ATOMIC, MODIFY, DELETE)

OBJECT\_HAS\_EXTERNAL\_LINKS\_PREVENTING\_DELETION (*version*)

OBJECT\_HAS\_INTERNAL\_LINKS\_PREVENTING\_DELETION (*version*)

OBJECT\_IS\_IN\_USE\_FOR\_DELETE (*version*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_HISTORY)

VERSION\_IS\_REQUIRED (*version*, COMPOSITE)

**9.4.4 VERSION\_REMOVE\_PREDECESSOR**

```
VERSION_REMOVE_PREDECESSOR (
    version      : Object_designator,
    predecessor : Object_designator
)
```

VERSION\_REMOVE\_PREDECESSOR removes the object *predecessor* as a predecessor of *version* in the graph of versions, by deleting the "predecessor" link from *version* to *predecessor* and its reverse "successor" link.

Write locks of the default mode are obtained on the deleted links.

**Errors**

ACCESS\_ERRORS (*predecessor*, ATOMIC, CHANGE, STABILIZE)

If *predecessor* is to be deleted:

ACCESS\_ERRORS (*predecessor* and its deleted components, ATOMIC, MODIFY, DELETE)

ACCESS\_ERRORS (*predecessor*, COMPOSITE, CHANGE)

ACCESS\_ERRORS (*version*, ATOMIC, MODIFY, WRITE\_LINKS)

If there is no "predecessor" link between *version* and *predecessor*:

LINK\_DOES\_NOT\_EXIST (*version*, "predecessor" link)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_HISTORY)



### 9.4.5 VERSION\_REVERSE

```

VERSION_REVERSE (
    version          : Object_designator,
    new_origin       : Object_designator,
    new_link         : Link_designator,
    on_same_volume_as : [ Object_designator ],
    access_mask      : Atomic_access_rights
)
    new_version      : Object_designator

```

VERSION\_REVERSE creates a new updatable version (a *revision*) of *version*. That is:

- A copy *new\_version* of *version* is created in the same way as

```
OBJECT_COPY (version, new_origin, new_link, nil, on_same_volume_as, access_mask)
```

where *reverse\_key* is not supplied.

- "predecessor" links with key 1 (one) are created from *new\_version* and each of its components to *version* and each of its corresponding components. As a consequence, *new\_version* and each of its components becomes a new successor of *version* and each of its corresponding components (i.e. a reverse "successor" link with a system-generated key is created).

Since the "predecessor" links are compositely stabilizing, *version* and its components are made stable.

*access\_mask* is used in conjunction with the default atomic ACL and default object owner of the calling process to specify the atomic ACL and the composite ACL of *new\_version*. *new\_version* and its components have the same integrity and confidentiality labels as the objects they have been copied from.

If *on\_same\_volume\_as* is supplied, the *new\_version* resides on the same volume as *on\_same\_volume\_as*. Otherwise, *new\_version* resides on the same volume as *version* and each component of *new\_version* resides on the same volume as the corresponding component of *version*.

If a replicated component is revised, its revision is not replicated.

Read locks of the default mode are obtained on all components of *version* and write locks of the default mode are obtained on the new links and components of *new\_version*.

A read lock of the default mode is obtained on *new\_origin* if the interpretation of *new\_link* implies the evaluation of any '+' or '++' key attribute values (see 23.1.2.7).

If *version* is an accounting log, its contents is preserved.

#### Errors

ACCESS\_ERRORS (*new\_origin*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*version*, COMPOSITE, CHANGE, APPEND\_IMPLICIT)

ACCESS\_ERRORS (*version*, COMPOSITE, READ, (READ\_LINKS, READ\_ATTRIBUTES, READ\_CONTENTS))

ACCESS\_ERRORS (*version*, COMPOSITE, CHANGE, STABILIZE)

If the reverse of *new\_link* is compositely stabilizing:

ACCESS\_ERRORS (*new\_origin*, COMPOSITE, CHANGE, STABILIZE)

ACCESS\_ERRORS (*on\_same\_volume\_as*, ATOMIC, SYSTEM\_ACCESS)

For each destination X of a duplicable external link L of a duplicated component:

ACCESS\_ERRORS (X, ATOMIC, CHANGE, APPEND\_IMPLICIT)

If L is atomically stabilizing:

ACCESS\_ERRORS (X, ATOMIC, CHANGE, STABILIZE)

If L is compositely stabilizing:

ACCESS\_ERRORS (X, COMPOSITE, CHANGE, STABILIZE)

CATEGORY\_IS\_BAD (*new\_origin*, *new\_link*, (COMPOSITION, EXISTENCE))

CATEGORY\_IS\_BAD (destination of *new\_link*, reverse of *new\_link*, IMPLICIT)

If *version* is a component of itself, and *new\_link* is a composition link:

COMPONENT\_ADDITION\_ERRORS (*new\_version*, *new\_link*)

CONTROL\_WOULD\_NOT\_BE\_GRANTED (*new\_version*)

EXTERNAL\_LINK\_IS\_BAD (*version*, COMPOSITE)

LINK\_EXISTS (*new\_origin*, *new\_link*)

OBJECT\_CANNOT\_BE\_STABILIZED (component of *version*)

OBJECT\_OWNER\_VALUE\_WOULD\_BE\_INCONSISTENT\_WITH\_ATOMIC\_ACL  
(*new\_object*)

If *version* is not a component of itself, *new\_link* is a composition link and *new\_origin* has  
OWNER granted or denied:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION (*new\_version*)

REFERENCE\_CANNOT\_BE\_ALLOCATED

TYPE\_OF\_OBJECT\_IS\_INVALID (*version*, COMPOSITE)

UPPER\_BOUND\_WOULD\_BE\_VIOLATED (*new\_origin*, *new\_link*)

USAGE\_MODE\_ON\_OBJECT\_TYPE\_WOULD\_BE\_VIOLATED ("object", type of *version*)

VALUE\_LIMIT\_ERRORS (*reverse\_key*)

VOLUME\_IS\_FULL (volume on which *on\_same\_volume\_as* resides)

#### 9.4.6 VERSION\_SNAPSHOT

```
VERSION_SNAPSHOT (
    version           : Object_designator,
    new_link_and_origin : [ Link_descriptor ],
    on_same_volume_as  : [ Object_designator ],
    access_mask       : Atomic_access_rights
)
    new_version      : Object_designator
```

VERSION\_SNAPSHOT creates a new stable version (a *snapshot*) of *version*. That is, if *new\_origin* and *new\_link* are the object designator and link designator respectively of *new\_link\_and\_origin*:

- A copy *new\_version* of *version* is created in the same way as

OBJECT\_COPY (*version*, *new\_origin*, *new\_link*, nil, *on\_same\_volume\_as*, *access\_mask*).

where *reverse\_link* is not supplied; except that if *new\_link\_and\_origin* is not supplied, then no new link of *new\_origin* is created

- The set of predecessors of each component of *new\_version* is the set of predecessors of its corresponding component of *version*. Then a "predecessor" link with key 1 (one) is created in such a way that each component of *new\_version* becomes the first predecessor of its corresponding component of *version*.

All the granted write and append discretionary access rights are suppressed (i.e. set to UNDEFINED) for all the components of *new\_version*.

The components of *version* are still updatable. *access\_mask* is used in conjunction with the default atomic ACL and default object owner of the calling process to specify the atomic ACL and the composite ACL of the created objects.

The components of *new\_version* are stabilized.

The created objects have the same integrity and confidentiality labels as the objects they have been copied from.

If *on\_same\_volume\_as* is supplied, the *new\_version* resides on the same volume as *on\_same\_volume\_as*. Otherwise, *new\_version* resides on the same volume as *version* and each component of *new\_version* resides on the same volume as the corresponding component of *version*.

If a component of *version* is replicated, its snapshot is not replicated.

The predecessor links are created even if their origins are stable.

Read locks of the default mode are obtained on all components of *version* to be copied and write locks of the default mode are obtained on the new links and components of *new\_version*. A read lock of the default mode is obtained on *new\_origin* if the interpretation of *new\_link* implies the evaluation of any '+' or '++' key attribute values (see 23.1.2.7).

If *version* is an accounting log, its contents is preserved.

## Errors

ACCESS\_ERRORS (*new\_origin*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*version*, COMPOSITE, MODIFY, APPEND\_LINKS)

If *version* or any component of *version* already has a predecessor:

ACCESS\_ERRORS (*version*, COMPOSITE, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (*version*, COMPOSITE, READ, (READ\_LINKS, READ\_ATTRIBUTES, READ\_CONTENTS))

ACCESS\_ERRORS (*new\_version*, COMPOSITE, CHANGE, STABILIZE)

If *version* has a predecessor, then for each predecessor X of each component of *version*:

ACCESS\_ERRORS (X, ATOMIC, CHANGE, (APPEND\_IMPLICIT, WRITE\_IMPLICIT))

If *new\_link* is provided and its reverse is compositely stabilizing:

ACCESS\_ERRORS (*new\_origin*, COMPOSITE, CHANGE, STABILIZE)

ACCESS\_ERRORS (*on\_same\_volume\_as*, ATOMIC, SYSTEM\_ACCESS)

For each destination X of a duplicable external link L of a duplicated component:

ACCESS\_ERRORS (X, ATOMIC, CHANGE, APPEND\_IMPLICIT)

If L is atomically stabilizing:  
ACCESS\_ERRORS (X, ATOMIC, CHANGE, STABILIZE)

If L is compositely stabilizing:  
ACCESS\_ERRORS (X, COMPOSITE, CHANGE, STABILIZE)

CATEGORY\_IS\_BAD (*new\_origin*, *new\_link*, (COMPOSITION, EXISTENCE, REFERENCE, DESIGNATION))

If *new\_link* has a reverse link:  
CATEGORY\_IS\_BAD (destination of *new\_link*, reverse of *new\_link*, IMPLICIT)

If *version* is a component of itself, and *new\_link* is a composition link:  
COMPONENT\_ADDITION\_ERRORS (*new\_version*, *new\_link*)

CONTROL\_WOULD\_NOT\_BE\_GRANTED (*new\_version*)

EXTERNAL\_LINK\_IS\_BAD (*version*, COMPOSITE)

REFERENCE\_CANNOT\_BE\_ALLOCATED

LINK\_EXISTS (*new\_origin*, *new\_link*)

OBJECT\_CANNOT\_BE\_STABILIZED (component of *version*)

OBJECT\_OWNER\_VALUE\_WOULD\_BE\_INCONSISTENT\_WITH\_ATOMIC\_ACL (*new\_object*)

If *version* is not a component of itself, *new\_link* is a composition link and *new\_origin* has OWNER granted or denied:  
OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION (*new\_version*)

TYPE\_OF\_OBJECT\_IS\_INVALID (*version*, COMPOSITE)

UPPER\_BOUND\_WOULD\_BE\_VIOLATED (*new\_origin*, *new\_link*)

USAGE\_MODE\_ON\_OBJECT\_TYPE\_WOULD\_BE\_VIOLATED ("object", type of *version*)

VALUE\_LIMIT\_ERRORS (*reverse\_key*)

VOLUME\_IS\_FULL (volume on which *on\_same\_volume\_as* resides)

#### 9.4.7 VERSION\_TEST\_ANCESTRY

```

VERSION_TEST_ANCESTRY (
    version1 : Object_designator,
    version2 : Object_designator
)
ancestry : Version_relation

```

VERSION\_TEST\_ANCESTRY tests the ancestry of the objects *version1* and *version2* in their graphs of versions. That is, it returns in *ancestry*:

- ANCESTOR\_VSN if there exists a series of "predecessor" links from *version2* to *version1*;
- DESCENDANT\_VSN if there exists a series of "predecessor" links from *version1* to *version2*;
- SAME\_VSN if *version1* is the same object as *version2*;

- RELATED\_VSN if there exist an object X which is neither *version1* nor *version2*, a series of "predecessor" links from *version1* to X, and a series of "predecessor" links from *version2* to X;
- UNRELATED\_VSN otherwise.

Read locks of the default mode are obtained on *version1* and on *version2* and on all the origins and destinations of the links in the series of links.

#### Errors

ACCESS\_ERRORS (*version1*, ATOMIC, READ, READ\_LINKS)

ACCESS\_ERRORS (element of version graph of *version1*, ATOMIC, READ, READ\_LINKS)

ACCESS\_ERRORS (*version2*, ATOMIC, READ, READ\_LINKS)

ACCESS\_ERRORS (element of version graph of *version2*, ATOMIC, READ, READ\_LINKS)

### 9.4.8 VERSION\_TEST\_DESCENT

```
VERSION_TEST_DESCENT (
    version1 : Object_designator,
    version2 : Object_designator
)
    descent : Version_relation
```

VERSION\_TEST\_DESCENT tests the descent of the objects *version1* and *version2* in their graphs of versions. That is, it returns:

- ANCESTOR\_VSN if there exists a series of "successor" links from *version1* to *version2*;
- DESCENDANT\_VSN if there exists a series of "successor" links from *version2* to *version1*;
- SAME\_VSN if *version1* is the same object as *version2*;
- RELATED\_VSN if there exist an object X which is neither *version1* nor *version2*, a series of "successor" links from *version1* to X, and a series of "successor" links from *version2* to X.
- UNRELATED\_VSN otherwise.

Read locks of the default mode are obtained on *version1* and on *version2* and on the all the origins and destinations of the links in the series of links.

#### Errors

ACCESS\_ERRORS (*version1*, ATOMIC, READ, READ\_LINKS)

ACCESS\_ERRORS (element of version graph of *version1*, ATOMIC, READ, READ\_LINKS)

ACCESS\_ERRORS (*version2*, ATOMIC, READ, READ\_LINKS)

ACCESS\_ERRORS (element of version graph of *version2*, ATOMIC, READ, READ\_LINKS)

## 10 Schema management

### 10.1 Schema management concepts

#### 10.1.1 Schema definition sets and the SDS directory

```

sds metasds:
  import object type system-object, system-process, system-common_root;
  import attribute type system-number, system-system_key;
  type_identifier: (read) string;
  sds_directory: child type of object with
  link
    known_sds: (navigate) non_duplicated existence link (sds_name: string) to sds;
    schemas_of: (navigate) implicit link to common_root reverse schemas;
  end sds_directory;
  sds: child type of object with
  link
    named_definition: (navigate) reference link (local_name: string) to type_in_sds
      reverse named_in_sds;
    in_working_schema_of: (navigate) non_duplicated designation link (number) to
      process;
  component
    definition: (navigate) exclusive composition link (type_identifier) to type_in_sds
      reverse in_sds;
  end sds;
  extend object type common_root with
  link
    schemas: (navigate) existence link to sds_directory reverse schemas_of;
  end common_root;
end metasds;

```

The SDS directory is an administrative object (see 9.1.2).

The "sds" components of the SDS directory represent the known SDSs of the PCTE installation (see 8.4):

- The "sds\_name" key of the "known\_sds" link from the SDS directory represents the SDS name of the SDS.
- The definition components of an SDS represent the types in SDS of the SDS (see 8.3). The key of the "named\_definition" link is the local name of the type in SDS. The destinations of the "named\_definition" links are a subset of the SDS object components. The destinations of the "in\_working\_schema" links are the known processes which are neither ready nor terminated and have included the SDS in their working schemas.

The destinations of the "in\_working\_schema" links are the known non-terminated processes which have included the SDS in their working schemas (see 8.5). An SDS and its types in SDS cannot be modified by using the operations defined in 10.2 while the SDS is included in the working schema of such a process.

NOTE - The predefined SDSs 'system', 'metasds', 'discretionary\_security', 'mandatory\_security', 'auditing', and 'accounting' are protected against modification, and so they cannot be extended directly. However, the predefined types can be imported into other SDSs and then extended, thus achieving the same effect. See 20.1.8.1.

### 10.1.2 Types

**sds** metasds:

type: **(protected) child type of object with attribute**

type\_identifier;

**link**

has\_type\_in\_sds: **(navigate) implicit link** (system\_key) to type\_in\_sds **reverse** of\_type;  
**end type**;

type\_in\_sds: **(protected) child type of object with attribute**

annotation: **string**;

creation\_or\_importation\_time: **(read) time**;

**link**

in\_sds: **(navigate) implicit link to sds reverse** definition;

of\_type: **(navigate) existence link to type reverse** has\_type\_in\_sds;

named\_in\_sds: **(navigate) implicit link to sds reverse** named\_definition;

**end type\_in\_sds**;

usage\_mode: **(read) natural**;

export\_mode: **(read) natural**;

maximum\_usage\_mode: **(read) natural**;

**end metasds**;

A "type" object represents a type (see 8.3):

- The "type\_identifier" attribute represents the type identifier.
- The destinations of the "has\_type\_in\_sds" links represent types in SDS associated with this type.

Further attribute types and link types are particular to the object types "object\_type" (see 10.1.3), "attribute\_type" (see 10.1.4), "link\_type" (see 10.1.5), and "enumerated\_type" (see 10.1.6).

A "type\_in\_sds" object represents a type in SDS (see 8.4):

- The destination of the "in\_sds" link represents the SDS to which the type in SDS belongs.
- The destination of the "named\_in\_sds" link represents the SDS in which the type in SDS has a local name. It is the same as the destination of the "in\_sds" link.
- The destination of the "of\_type" link represents the type associated with the type in SDS.
- The usage mode, export mode, and maximum usage mode represent the definition modes of the type in SDS. A set of definition mode values is represented as the sum of the powers of 2 representing its elements as follows:

- . CREATE : 1
- . DELETE : 2
- . READ : 4
- . WRITE : 8
- . NAVIGATE : 16



- The annotation is the complete name of the type when it is created, but may be changed by the user.
- The creation or importation time is the system time when the type in SDS was created or imported into the SDS.

Further attribute types and link types are particular to the object types "object\_type\_in\_sds" (see 10.1.3), "attribute\_type\_in\_sds" (see 10.1.4), "link\_type\_in\_sds" (see 10.1.5), and "enumeration\_type\_in\_sds" (see 10.1.6).

### 10.1.3 Object types

**sds metasds:**

**object\_type:** (**protected**) **child type of type with attribute**

**contents\_type:** (**read**) **enumeration** (FILE\_TYPE, PIPE\_TYPE, DEVICE\_TYPE, AUDIT\_FILE\_TYPE, ACCOUNTING\_LOG\_TYPE, NO\_CONTENTS\_TYPE) := NO\_CONTENTS\_TYPE;

**link**

**parent\_type:** (**navigate**) **reference link** (number) **to** object\_type **reverse** child\_type;  
**child\_type:** (**navigate**) **implicit link** (system\_key) **to** object\_type **reverse** parent\_type;

**end object\_type;**

**object\_type\_in\_sds:** (**protected**) **child type of type\_in\_sds with attribute**

**usage\_mode;**  
**export\_mode;**  
**maximum\_usage\_mode;**

**link**

**in\_attribute\_set:** (**navigate**) **reference link** (number) **to** attribute\_type\_in\_sds **reverse** is\_attribute\_of;  
**in\_link\_set:** (**navigate**) **reference link** (number) **to** link\_type\_in\_sds **reverse** is\_link\_of;  
**is\_destination\_of:** (**navigate**) **reference link** (number) **to** link\_type\_in\_sds **reverse** in\_destination\_set;

**end object\_type\_in\_sds;**

**end metasds;**

An "object\_type" object represents an object type (see 8.3.1):

- The destinations of the "parent\_type" links represent the parent types of the object type.
- The destinations of the "child\_type" links represent the child types of the object type.

An "object\_type\_in\_sds" object represents an object type in SDS (see 8.4.1):

- The destinations of the "in\_attribute\_set" links represent the direct attribute types in SDS of the object type in SDS.
- The destinations of the "in\_link\_set" links represent the direct outgoing link types in SDS of the object type in SDS. The component object types of the object type in SDS are represented by the destinations of the "in\_destination\_set" links of the destinations of the "in\_link\_set" links with category COMPOSITION.
- The destinations of the "is\_destination\_of" links represent the link types in SDS of which this object type is a destination object type.

#### 10.1.4 Attribute types

**sds** metasds:

duplication: (**read**) **enumeration** (DUPLICATED, NON\_DUPLICATED):= DUPLICATED;

key\_attribute\_of: (**navigate**) **implicit link** (system\_key) **to** link\_type **reverse** key\_attribute;

attribute\_type: (**protected**) **child type of type with attribute**

duplication;

**end** attribute\_type;

string\_attribute\_type: (**protected**) **child type of attribute\_type with attribute**

string\_initial\_value: (**read**) **string**;

**link**

key\_attribute\_of;

**end** string\_attribute\_type;

integer\_attribute\_type: (**protected**) **child type of attribute\_type with attribute**

integer\_initial\_value: (**read**) **integer**;

**end** integer\_attribute\_type;

natural\_attribute\_type: (**protected**) **child type of attribute\_type with attribute**

natural\_initial\_value: (**read**) **natural**;

**link**

key\_attribute\_of;

**end** natural\_attribute\_type;

float\_attribute\_type: (**protected**) **child type of attribute\_type with attribute**

float\_initial\_value: (**read**) **float**;

**end** float\_attribute\_type;

boolean\_attribute\_type: (**protected**) **child type of attribute\_type with attribute**

boolean\_initial\_value: (**read**) **boolean**;

**end** boolean\_attribute\_type;

time\_attribute\_type: (**protected**) **child type of attribute\_type with attribute**

time\_initial\_value: (**read**) **time**;

**end** time\_attribute\_type;

enumeration\_attribute\_type: (**protected**) **child type of attribute\_type with attribute**

initial\_value\_position: (**read**) **natural**;

**component**

enumerals: (**navigate**) **composition link** [1 .. ] (position: **natural**) **to** enumerals\_type **reverse** enumerals\_of;

**end** enumeration\_attribute\_type;

```

attribute_type_in_sds: (protected) child type of type_in_sds with
attribute
    usage_mode;
    export_mode;
    maximum_usage_mode;
link
    is_attribute_of: (navigate) reference link (number) to object_type_in_sds,
    link_type_in_sds reverse in_attribute_set;
end attribute_type_in_sds;
end metasds;

```

"Attribute\_type" objects represent attribute types (see 8.3.2). They are divided into child types according to value type.

- The initial value attribute represents the initial value of the attribute type. For string, integer, natural, float, boolean, and time attribute types, it is an actual value of the value type. For an enumeration attribute type, it is a non-negative integer defining the position of the initial value within the enumeration type.
- For a natural or a string attribute type, the destinations of the "key\_attribute\_of" links represent the link types for which this attribute type is a key attribute type.
- For an enumeration attribute type, the destinations of the "enumerals" links represent the enumerals types of the value type. The "position" key attribute represents the ordering of the enumerals types: it must take successive values 0, 1, 2, 3, ... .

An "attribute\_type\_in\_sds" object represents an attribute type in SDS (see 8.4.2):

- The destinations of the "in\_attribute\_set" links represent the object types in SDS and link types in SDS for which the attribute type is a direct attribute type.

### 10.1.5 Link types

sds metasds:

```

link_type : (protected) child type of type with
attribute

```

```

    category: (read) enumeration (COMPOSITION, EXISTENCE, REFERENCE, IMPLICIT,
    DESIGNATION) := COMPOSITION;
    lower_bound: (read) natural := 0;
    upper_bound: (read) natural := MAX_NATURAL_ATTRIBUTE;
    stability: (read) enumeration (ATOMIC_STABLE, COMPOSITE_STABLE, NON_STABLE)
    := NON_STABLE;
    exclusiveness: (read) enumeration (SHARABLE, EXCLUSIVE) := SHARABLE;
    duplication;

```

**link**

```

    reverse: (navigate) reference link to link_type;
    key_attribute: (navigate) reference link (key_number: natural) to string_attribute_type,
    natural_attribute_type reverse key_attribute_of;
end link_type;

```

link\_type\_in\_sds: **child type of** type\_in\_sds **with attribute**

usage\_mode;  
export\_mode;  
maximum\_usage\_mode;

**link**

in\_attribute\_set;  
is\_link\_of: **(navigate) reference link** (number) **to** object\_type\_in\_sds **reverse**  
in\_link\_set;  
in\_destination\_set: **(navigate) reference link** (number) **to** object\_type\_in\_sds **reverse**  
is\_destination\_of;

**end** link\_type\_in\_sds;

**end** metasds;

A "link\_type" object represents a link type (see 8.3.3):

- The "category" attribute represents the category of the link type.
- The "lower\_bound" and "upper\_bound" attributes represent the lower bound and upper bound, respectively, of the link type. For MAX\_NATURAL\_ATTRIBUTE see 24.1.
- The "stability" attribute represents the stability of the link type.
- The "exclusiveness" attribute represents the exclusiveness of the link type.
- The "duplication" attribute represents the duplication property of the link type.
- The destination of the "reverse" link represents the reverse link type of the link type.
- The destinations of the "key\_attribute" links represent the key attribute types of the link type.

A "link\_type\_in\_sds" object represents a link type in SDS (see 8.4.3):

- The destinations of the "in\_attribute\_set" links represent the non-key attribute types of the link type in SDS.
- The destinations of the "is\_link\_of" links represent the origin object types in SDS of the link type in SDS.
- The destinations of the "in\_destination\_set" links represent the destination object types in SDS of the link type in SDS.

### 10.1.6 Enumeral types

**sds** metasds:

enumeration\_type: **(protected) child type of** type **with link**

enumeration\_of: **(navigate) implicit link** (system\_key) **to** enumeration\_attribute\_type  
**reverse** enumeration;

**end** enumeration\_type;

enumeration\_type\_in\_sds: **(protected) child type of** type\_in\_sds **with attribute**

image: **(read) string**;

**end** enumeration\_type\_in\_sds;

**end** metasds;

An "enumeration\_type" object represents an enumeration type (see 8.3.4).

The destinations of the "enumeration\_of" links represent the attribute types of which the enumeration type is a possible value.

An "enumeration\_type\_in\_sds" object represents an enumeration type in SDS (see 8.4.4); the "image" attribute represents the image of the enumeration type in SDS.

### 10.1.7 Datatypes for schema management

Enumeration\_values = **seq1 of** Enumeration\_type\_nominator\_in\_sds

Key\_types\_in\_sds = **seq of** Attribute\_type\_nominator\_in\_sds

Attribute\_scan\_kind = OBJECT | OBJECT\_ALL | LINK\_KEY | LINK\_NON\_KEY

Link\_scan\_kind = ORIGIN | ORIGIN\_ALL | DESTINATION | DESTINATION\_ALL | KEY | NON\_KEY

Object\_scan\_kind = CHILD | DESCENDANT | PARENT | ANCESTOR | ATTRIBUTE | ATTRIBUTE\_ALL | LINK\_ORIGIN | LINK\_ORIGIN\_ALL | LINK\_DESTINATION | LINK\_DESTINATION\_ALL

These datatypes are used as parameter and result types of operations defined in 10.2, 10.3, and 10.4.

## 10.2 SDS update operations

### 10.2.1 SDS\_ADD\_DESTINATION

```
SDS_ADD_DESTINATION (
    sds          : Sds_designator,
    link_type    : Link_type_nominator_in_sds,
    object_type  : Object_type_nominator_in_sds
)
```

SDS\_ADD\_DESTINATION extends the set of destination object types of the link type in SDS *link\_type\_in\_sds* associated with the link type *link\_type* in the SDS *sds* to include the object type in SDS *object\_type\_in\_sds* associated with the object type *object\_type* in *sds*.

If *link\_type* has a reverse link type *reverse*, then *reverse* is applied to *object\_type*.

An "in\_destination\_set" link from *link\_type\_in\_sds* to *object\_type\_in\_sds* and its reverse "is\_destination\_of" link are created.

If *link\_type* has a reverse link type *reverse* then an "in\_link\_set" link from *object\_type\_in\_sds* to the "link\_type\_in\_sds" object *reverse\_link\_type\_in\_sds* associated with *reverse* in *sds*, and its reverse "is\_link\_of" link, are created.

Write locks of the default mode are obtained on the created links.

#### Errors

ACCESS\_ERRORS (*object\_type\_in\_sds*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*link\_type\_in\_sds*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*sds*, ATOMIC, READ, NAVIGATE)

ACCESS\_ERRORS (*reverse\_link\_type\_in\_sds*, ATOMIC, MODIFY, APPEND\_LINKS)

OBJECT\_TYPE\_IS\_ALREADY\_IN\_DESTINATION\_SET (*link\_type*, *object\_type*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)  
 SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)  
 SDS\_IS\_PREDEFINED (*sds*)  
 SDS\_IS\_UNKNOWN (*sds*)  
 TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *link\_type*)  
 TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *object\_type*)

### 10.2.2 SDS\_APPLY\_ATTRIBUTE\_TYPE

```

SDS_APPLY_ATTRIBUTE_TYPE (
    sds           : Sds_designator,
    attribute_type : Attribute_type_nominator_in_sds,
    type          : Object_type_nominator_in_sds | Link_type_nominator_in_sds
)
  
```

SDS\_APPLY\_ATTRIBUTE\_TYPE extends the object type or link type *type* by the application of the attribute type *attribute\_type* in the SDS *sds*.

An "in\_attribute\_set" link and its reverse "is\_attribute\_of" link are created between the type in SDS *type\_in\_sds* associated with *type* in *sds* and the attribute type in SDS *attribute\_type\_in\_sds* associated with *attribute\_type* in *sds*.

Write locks of the default mode are obtained on the created links.

#### Errors

ACCESS\_ERRORS (*type\_in\_sds*, ATOMIC, MODIFY, APPEND\_LINKS)  
 ACCESS\_ERRORS (*attribute\_type\_in\_sds*, ATOMIC, MODIFY, APPEND\_LINKS)  
 ACCESS\_ERRORS (*sds*, ATOMIC, READ, NAVIGATE)  
 KEY\_ATTRIBUTE\_TYPE\_APPLY\_IS\_FORBIDDEN (*attribute\_type*)  
 LINK\_TYPE\_CATEGORY\_IS\_BAD (*link\_type*, (COMPOSITION, EXISTENCE, REFERENCE, DESIGNATION))  
 PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)  
 SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)  
 SDS\_IS\_PREDEFINED (*sds*)  
 SDS\_IS\_UNKNOWN (*sds*)  
 TYPE\_CANNOT\_BE\_APPLIED\_TO\_LINK\_TYPE (*type*, *attribute\_type*)  
 TYPE\_IS\_ALREADY\_APPLIED (*sds*, *attribute\_type*, *type*)  
 TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *attribute\_type*)  
 TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *type*)

### 10.2.3 SDS\_APPLY\_LINK\_TYPE

```

SDS_APPLY_LINK_TYPE (
    sds           : Sds_designator,
    link_type      : Link_type_nominator_in_sds,
    object_type    : Object_type_nominator_in_sds
)
  
```

SDS\_APPLY\_LINK\_TYPE extends the object type *object\_type* in the SDS *sds* by the application of the link type *link\_type*. If *link\_type* has a reverse link type *reverse* then the destination set of *reverse* is extended to include the object type *object\_type*.

An "in\_link\_set" link from the object type in SDS *object\_type\_in\_sds* associated with *object\_type* in *sds* to the link type in SDS *link\_type\_in\_sds* associated with *link\_type* in *sds*, and its reverse "is\_link\_of" link, are created. If *link\_type* has a reverse link type *reverse*, an "in\_destination\_set" link from the link type in SDS associated with *reverse* in *sds* to *object\_type\_in\_sds*, and its reverse "is\_destination\_of" link, are created.

Write locks of the default mode are obtained on the created links.

### Errors

ACCESS\_ERRORS (*link\_type\_in\_sds*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*object\_type\_in\_sds*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*sds*, ATOMIC, READ, NAVIGATE)

ACCESS\_ERRORS (reverse of *link\_type\_in\_sds*, ATOMIC, MODIFY, APPEND\_LINKS)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)

SDS\_IS\_PREDEFINED (*sds*)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_ALREADY\_APPLIED (*sds*, *link\_type*, *object\_type*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *object\_type*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *link\_type*)

## 10.2.4 SDS\_CREATE\_BOOLEAN\_ATTRIBUTE\_TYPE

```
SDS_CREATE_BOOLEAN_ATTRIBUTE_TYPE (
    sds                : Sds_designator,
    local_name         : [ Name ],
    initial_value      : [ Boolean ],
    duplication        : Duplication
)
new_type             : Attribute_type_nominator_in_sds
```

SDS\_CREATE\_BOOLEAN\_ATTRIBUTE\_TYPE creates a new boolean attribute type *new\_type*, and its associated attribute type in SDS *new\_type\_in\_sds* in the SDS *sds*.

The operation creates a "definition" link from *sds* to *new\_type\_in\_sds*; the key of the link is the system-assigned type identifier of *new\_type*. The operation also creates an "of\_type" link from *new\_type\_in\_sds* to *new\_type*.

If *local\_name* is supplied, a "named\_definition" link is created from *sds* to *new\_type\_in\_sds* with *local\_name* as key, together with its reverse "named\_in\_sds" link.

The type identifier of *new\_type* is set to an implementation-defined value which identifies the type within the PCTE installation. The duplication of *new\_type* is set to *duplication*. The boolean initial value of *new\_type* is set to *initial\_value* if supplied, and otherwise to **false**.



The three definition mode attributes of *new\_type\_in\_sds* are set to 12, representing READ\_MODE and WRITE\_MODE, and its creation or importation time is set to the system time. If *local\_name* is supplied, the annotation of *new\_type\_in\_sds* is set to the complete name of the created type; otherwise it is set to the empty string.

The new objects reside on the same volume as *sds*. Their access control lists are built using the default atomic ACL and the default object owner of the calling process, and their confidentiality labels and integrity labels are set to be equal to the current confidentiality context and integrity context, respectively, of the calling process.

For each created object, an "object\_on\_volume" link is created from the volume on which the object resides to the object. The key of the link is the exact identifier of the object.

Write locks of the default mode are obtained on the created objects and links except the new "object\_on\_volume" links.

### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, MODIFY, APPEND\_LINKS)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_DEFINITION\_NAME\_SIZE)

If *sds* has OWNER granted or denied:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION  
(*new\_type\_in\_sds*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)

SDS\_IS\_PREDEFINED (*sds*)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*sds*, *local\_name*)

TYPE\_NAME\_IS\_INVALID (*local\_name*)

### 10.2.5 SDS\_CREATE\_DESIGNATION\_LINK\_TYPE

```
SDS_CREATE_DESIGNATION_LINK_TYPE (
    sds           : Sds_designator,
    local_name    : [ Name ],
    lower_bound   : [ Natural ],
    upper_bound   : [ Natural ],
    duplication    : Duplication,
    key_types     : Key_types_in_sds
)
    new_type      : Link_type_nominator_in_sds
```

SDS\_CREATE\_DESIGNATION\_LINK\_TYPE creates a new designation link type *new\_type* and its associated link type in SDS *new\_type\_in\_sds* in the SDS *sds*.

The operation creates a "definition" link from *sds* to *new\_type\_in\_sds*; the key of the link is the system-assigned type identifier of *new\_type*. The operation also creates an "of\_type" link from *new\_type\_in\_sds* to *new\_type*.

If *local\_name* is supplied, a "named\_definition" link is created from *sds* to *new\_type\_in\_sds* with *local\_name* as key, together with its reverse "named\_in\_sds" link.

The type identifier of *new\_type* is set to an implementation-defined value which identifies the type within the PCTE installation.

The three definition mode attributes of *new\_type\_in\_sds* are set to 19, representing CREATE\_MODE, DELETE\_MODE, and NAVIGATE\_MODE, and its creation or importation time is set to the system time. If *local\_name* is supplied, the annotation of *new\_type\_in\_sds* is set to the complete name of the created type; otherwise it is set to the empty string.

The lower bound, upper bound (if provided), and duplication of *new\_type* are set from the parameters of the same names. If *new\_type* is of cardinality many, for each attribute type of *key\_types*, a "key\_attribute" link is created from *new\_type* to that attribute type. The keys of these links correspond to the order of the key attribute types in *key\_types* starting at 0 and incremented by 1. The category of *new\_type* is set to DESIGNATION.

The sets of origin object types in SDS and destination object types in SDS of *new\_type* are initially empty.

The new objects reside on the same volume as *sds*. Their access control lists are built using the default atomic ACL and the default object owner of the calling process, and their confidentiality labels and integrity labels are set to be equal to the current confidentiality context and integrity context, respectively, of the calling process.

For each created object, an "object\_on\_volume" link is created from the volume on which the object resides to the object. The key of the link is the exact identifier of the object.

Write locks of the default mode are obtained on the created objects and links except the new "object\_on\_volume" links.

### Errors

ACCESS\_ERRORS (element of *key\_types*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

ACCESS\_ERRORS (*sds*, ATOMIC, MODIFY, APPEND\_LINKS)

KEY\_TYPE\_IS\_BAD (element of *key\_types*)

KEY\_TYPES\_ARE\_MULTIPLE (*key\_types*)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_DEFINITION\_NAME\_SIZE)

LINK\_TYPE\_PROPERTIES\_ARE\_INCONSISTENT (DESIGNATION, *lower\_bound*, *upper\_bound*, SHARABLE, NON\_STABLE, *duplication*)

LINK\_TYPE\_PROPERTIES\_AND\_KEY\_TYPES\_ARE\_INCONSISTENT (DESIGNATION, *lower\_bound*, *upper\_bound*, SHARABLE, NON\_STABLE, *duplication*, *key\_types*)

If *sds* has OWNER granted or denied:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION  
(*new\_type\_in\_sds*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)

SDS\_IS\_PREDEFINED (*sds*)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, element of *key\_types*)

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*sds*, *local\_name*)

TYPE\_NAME\_IS\_INVALID (*local\_name*)

### 10.2.6 SDS\_CREATE\_ENUMERAL\_TYPE

```

SDS_CREATE_ENUMERAL_TYPE (
    sds          : Sds_designator,
    local_name   : [ Name ]
)
    new_type     : Enumeral_type_nominator_in_sds

```

SDS\_CREATE\_ENUMERAL\_TYPE creates a new enumerals type *new\_type* and its associated enumerals type in SDS *new\_type\_in\_sds* in the SDS *sds*.

The operation creates a "definition" link from *sds* to *new\_type\_in\_sds*; the key of the link is the system-assigned type identifier of *new\_type*. The operation also creates an "of\_type" link from *new\_type\_in\_sds* to *new\_type*.

If *local\_name* is supplied, a "named\_definition" link is created from *sds* to *new\_type\_in\_sds* with *local\_name* as key, together with its reverse "named\_in\_sds" link.

The type identifier of *new\_type* is set to an implementation-defined value which identifies the type within the PCTE installation.

The creation or importation time of *new\_type\_in\_sds* is set to the system time. If *local\_name* is supplied, the annotation of *new\_type\_in\_sds* is set to the complete name of the created type; otherwise it is set to the empty string. The image of *new\_type\_in\_sds* is set to the empty string.

The new objects reside on the same volume as *sds*. Their access control lists are built using the default atomic ACL and the default object owner of the calling process, and their confidentiality labels and integrity labels are set to be equal to the current confidentiality context and integrity context, respectively, of the calling process.

For each created object, an "object\_on\_volume" link is created from the volume on which the object resides to the object. The key of the link is the exact identifier of the object.

Write locks of the default mode are obtained on the created objects and links except the new "object\_on\_volume" links.

#### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, MODIFY, APPEND\_LINKS)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_DEFINITION\_NAME\_SIZE)

If *sds* has OWNER granted or denied:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION  
(*new\_type\_in\_sds*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)

SDS\_IS\_PREDEFINED (*sds*)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*sds*, *local\_name*)

TYPE\_NAME\_IS\_INVALID (*local\_name*)

### 10.2.7 SDS\_CREATE\_ENUMERATION\_ATTRIBUTE\_TYPE

```

SDS_CREATE_ENUMERATION_ATTRIBUTE_TYPE (
    sds           : Sds_designator,
    local_name    : [ Name ],
    values        : Enumeration_values,
    duplication   : Duplication,
    initial_value : [ Natural ]
)
    new_type      : Attribute_type_nominator_in_sds

```

SDS\_CREATE\_ENUMERATION\_ATTRIBUTE\_TYPE creates a new enumeration attribute type *new\_type*, and its associated attribute type in SDS *new\_type\_in\_sds* in the SDS *sds*.

The operation creates a "definition" link from *sds* to *new\_type\_in\_sds*; the key of the link is the system-assigned type identifier of *new\_type*. The operation also creates an "of\_type" link from *new\_type\_in\_sds* to *new\_type*.

"enumerals" links are created from *new\_type* to the enumerals specified in *values*. The "position" key attributes of the links are set according to the sequential order of the type nominators in *value*, starting at 0 and incremented by 1 for each enumeral type. The reverse "enumerals\_of" links are also created.

If *local\_name* is supplied, a "named\_definition" link is created from *sds* to *new\_type\_in\_sds* with *local\_name* as key, together with its reverse "named\_in\_sds" link.

The type identifier of *new\_type* is set to an implementation-defined value which identifies the type within the PCTE installation. The duplication of *new\_type* is set to *duplication*. The initial value position of *new\_type* is set to *initial\_value* if supplied, and otherwise to 0.

The three definition mode attributes of *new\_type\_in\_sds* are set to 12, representing READ\_MODE and WRITE\_MODE, and its creation or importation time is set to the system time. If *local\_name* is supplied, the annotation of *new\_type\_in\_sds* is set to the complete name of the created type; otherwise it is set to the empty string.

The new objects reside on the same volume as *sds*. Their access control lists are built using the default atomic ACL and the default object owner of the calling process, and their confidentiality labels and integrity labels are set to be equal to the current confidentiality context and integrity context, respectively, of the calling process.

For each created object, an "object\_on\_volume" link is created from the volume on which the object resides to the object. The key of the link is the exact identifier of the object.

Write locks of the default mode are obtained on the created objects and links except the new "object\_on\_volume" links.

A write lock of the default mode is obtained on any enumeral type specified in *values* if the OWNER discretionary access right is granted for a group to *new\_type* (the default object owner group), and a different OWNER discretionary access right exists for the same group to that enumeral type.

#### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (element of *values*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

For each enumerational type in SDS *E* associated with an element of *values*:

ACCESS\_ERRORS (*E*, ATOMIC, READ, READ\_ATTRIBUTES)  
 COMPONENT\_ADDITION\_ERRORS (enumerational type, "enumerational" link)  
 ENUMERATIONAL\_TYPES\_ARE\_MULTIPLE (*values*)  
 ENUMERATION\_ATTRIBUTE\_WOULD\_HAVE\_NO\_ENUMERATIONAL\_TYPES (*values*)  
 ENUMERATION\_VALUE\_IS\_OUT\_OF\_RANGE (*initial\_value*, *values*)  
 IMAGE\_IS\_DUPLICATED (*values*, *sds*)  
 LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_DEFINITION\_NAME\_SIZE)

If *sds* has OWNER granted or denied:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION  
 (*new\_type\_in\_sds*)  
 PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)  
 SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)  
 SDS\_IS\_PREDEFINED (*sds*)  
 SDS\_IS\_UNKNOWN (*sds*)  
 TYPE\_IS\_UNKNOWN\_IN\_SDS(*sds*, element of *values*)  
 TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*sds*, *local\_name*)  
 TYPE\_NAME\_IS\_INVALID (*local\_name*)

### 10.2.8 SDS\_CREATE\_FLOAT\_ATTRIBUTE\_TYPE

```
SDS_CREATE_FLOAT_ATTRIBUTE_TYPE (
  sds           : Sds_designator,
  local_name    : [ Name ],
  initial_value : [ Float ],
  duplication   : Duplication
)
  new_type      : Attribute_type_nominator_in_sds
```

SDS\_CREATE\_FLOAT\_ATTRIBUTE\_TYPE creates a new float attribute type *new\_type*, and its associated attribute type in SDS *new\_type\_in\_sds* in the SDS *sds*.

The operation creates a "definition" link from *sds* to *new\_type\_in\_sds*; the key of the link is the system-assigned type identifier of *new\_type*. The operation also creates an "of\_type" link from *new\_type\_in\_sds* to *new\_type*.

If *local\_name* is supplied, a "named\_definition" link is created from *sds* to *new\_type\_in\_sds* with *local\_name* as key, together with its reverse "named\_in\_sds" link.

The type identifier of *new\_type* is set to an implementation-defined value which identifies the type within the PCTE installation. The duplication of *new\_type* is set to *duplication*. The float initial value of *new\_type* is set to *initial\_value* if supplied, and otherwise to 0.0 (zero).

The three definition mode attributes of *new\_type\_in\_sds* are set to 12, representing READ\_MODE and WRITE\_MODE, and its creation or importation time is set to the system time. If *local\_name* is supplied, the annotation of *new\_type\_in\_sds* is set to the complete name of the created type; otherwise it is set to the empty string.

The new objects reside on the same volume as *sds*. Their access control lists are built using the default atomic ACL and the default object owner of the calling process, and their confidentiality labels and integrity labels are set to be equal to the current confidentiality context and integrity context, respectively, of the calling process.

For each created object, an "object\_on\_volume" link is created from the volume on which the object resides to the object. The key of the link is the exact identifier of the object.

Write locks of the default mode are obtained on the created objects and links except the new "object\_on\_volume" links.

### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, MODIFY, APPEND\_LINKS)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_DEFINITION\_NAME\_SIZE)

If *sds* has OWNER granted or denied:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION  
(*new\_type\_in\_sds*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)

SDS\_IS\_PREDEFINED (*sds*)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*sds*, *local\_name*)

TYPE\_NAME\_IS\_INVALID (*local\_name*)

VALUE\_LIMIT\_ERRORS (*initial\_value*)

### 10.2.9 SDS\_CREATE\_INTEGER\_ATTRIBUTE\_TYPE

```
SDS_CREATE_INTEGER_ATTRIBUTE_TYPE (
    sds                Sds_designator,
    local_name         : [ Name ],
    initial_value      : [ Integer ],
    duplication        : Duplication
)
    new_type          : Attribute_type_nominator_in_sds
```

SDS\_CREATE\_INTEGER\_ATTRIBUTE\_TYPE creates a new integer attribute type *new\_type*, and its associated attribute type in SDS *new\_type\_in\_sds* in the SDS *sds*.

The operation creates a "definition" link from *sds* to *new\_type\_in\_sds*; the key of the link is the system-assigned type identifier of *new\_type*. The operation also creates an "of\_type" link from *new\_type\_in\_sds* to *new\_type*.

If *local\_name* is supplied, a "named\_definition" link is created from *sds* to *new\_type\_in\_sds* with *local\_name* as key, together with its reverse "named\_in\_sds" link.

The type identifier of *new\_type* is set to an implementation-defined value which identifies the type within the PCTE installation. The duplication of *new\_type* is set to *duplication*. The integer initial value of *new\_type* is set to *initial\_value* if supplied, and otherwise to 0 (zero).

The three definition mode attributes of *new\_type\_in\_sds* are set to 12, representing READ\_MODE and WRITE\_MODE, and its creation or importation time is set to the system

time. If *local\_name* is supplied, the annotation of *new\_type\_in\_sds* is set to the complete name of the created type; otherwise it is set to the empty string.

The new objects reside on the same volume as *sds*. Their access control lists are built using the default atomic ACL and the default object owner of the calling process, and their confidentiality labels and integrity labels are set to be equal to the current confidentiality context and integrity context, respectively, of the calling process.

For each created object, an "object\_on\_volume" link is created from the volume on which the object resides to the object. The key of the link is the exact identifier of the object.

Write locks of the default mode are obtained on the created objects and links except the new "object\_on\_volume" links.

### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, MODIFY, APPEND\_LINKS)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_DEFINITION\_NAME\_SIZE)

If *sds* has OWNER granted or denied:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION  
(*new\_type\_in\_sds*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)

SDS\_IS\_UNKNOWN (*sds*)

SDS\_IS\_PREDEFINED (*sds*)

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*sds*, *local\_name*)

TYPE\_NAME\_IS\_INVALID (*local\_name*)

VALUE\_LIMIT\_ERRORS (*initial\_value*)

### 10.2.10 SDS\_CREATE\_NATURAL\_ATTRIBUTE\_TYPE

```
SDS_CREATE_NATURAL_ATTRIBUTE_TYPE (
    sds           : Sds_designator,
    local_name    : [ Name ],
    initial_value : [ Natural ],
    duplication   : Duplication
)
new_type       : Attribute_type_nominator_in_sds
```

SDS\_CREATE\_NATURAL\_ATTRIBUTE\_TYPE creates a new natural attribute type *new\_type*, and its associated attribute type in SDS *new\_type\_in\_sds* in the SDS *sds*.

The operation creates a "definition" link from *sds* to *new\_type\_in\_sds*; the key of the link is the system-assigned type identifier of *new\_type*. The operation also creates an "of\_type" link from *new\_type\_in\_sds* to *new\_type*.

If *local\_name* is supplied, a "named\_definition" link is created from *sds* to *new\_type\_in\_sds* with *local\_name* as key, together with its reverse "named\_in\_sds" link.

The type identifier of *new\_type* is set to an implementation-defined value which identifies the type within the PCTE installation. The duplication of *new\_type* is set to *duplication*. The natural initial value of *new\_type* is set to *initial\_value* if supplied, and otherwise to 0 (zero).



The three definition mode attributes of *new\_type\_in\_sds* are set to 12, representing READ\_MODE and WRITE\_MODE, and its creation or importation time is set to the system time. If *local\_name* is supplied, the annotation of *new\_type\_in\_sds* is set to the complete name of the created type; otherwise it is set to the empty string.

The new objects reside on the same volume as *sds*. Their access control lists are built using the default atomic ACL and the default object owner of the calling process, and their confidentiality labels and integrity labels are set to be equal to the current confidentiality context and integrity context, respectively, of the calling process.

For each created object, an "object\_on\_volume" link is created from the volume on which the object resides to the object. The key of the link is the exact identifier of the object.

Write locks of the default mode are obtained on the created objects and links except the new "object\_on\_volume" links.

### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, MODIFY, APPEND\_LINKS)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_DEFINITION\_NAME\_SIZE)

If *sds* has OWNER granted or denied:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION  
(*new\_type\_in\_sds*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)

SDS\_IS\_PREDEFINED (*sds*)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*sds*, *local\_name*)

TYPE\_NAME\_IS\_INVALID (*local\_name*)

VALUE\_LIMIT\_ERRORS (*initial\_value*)

### 10.2.11 SDS\_CREATE\_OBJECT\_TYPE

```
SDS_CREATE_OBJECT_TYPE (
    sds           : Sds_designator,
    local_name    : [ Name ],
    parents       : Object_type_nominators_in_sds
)
    new_type     : Object_type_nominator_in_sds
```

SDS\_CREATE\_OBJECT\_TYPE creates a new object type *new\_type* and its associated object type in SDS *new\_type\_in\_sds* in the SDS *sds*.

The contents type of *new\_type* is determined as follows. If one or more of the object types in *parents* has a contents type (if more than one, they must all have the same contents type) then *new\_type* has that contents type; otherwise *new\_type* has no contents type.

The type identifier of *new\_type* is set to an implementation-defined value which identifies the type within the PCTE installation.

The operation creates a "definition" link from *sds* to *new\_type\_in\_sds*; the key of the link is the system-assigned type identifier of *new\_type*. The operation also creates an "of\_type" link from *new\_type\_in\_sds* to *new\_type*.

If *local\_name* is supplied, a "named\_definition" link is created from *sds* to *new\_type\_in\_sds* with *local\_name* as key, together with its reverse "named\_in\_sds" link. "parent\_type" links are created from *new\_type* to each of *parents*, together with their reverse "child\_type" links.

The three definition mode attributes of *new\_type\_in\_sds* are set to 1, representing CREATE\_MODE, and its creation or importation time is set to the system time. If *local\_name* is supplied, the annotation of *new\_type\_in\_sds* is set to the complete name of the created type; otherwise it is set to the empty string.

The new objects reside on the same volume as *sds*. Their access control lists are built using the default atomic ACL and the default object owner of the calling process and their confidentiality labels and integrity labels are set to be equal to the current confidentiality context and integrity context, respectively, of the calling process.

For each created object, an "object\_on\_volume" link is created from the volume on which the object resides to the object. The key of the link is the exact identifier of the object.

Write locks of the default mode are obtained on the created objects and links except the new "object\_on\_volume" links.

### Errors

ACCESS\_ERRORS (element of *parents*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

ACCESS\_ERRORS (*sds*, ATOMIC, MODIFY, APPEND\_LINKS)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_DEFINITION\_NAME\_SIZE)

OBJECT\_TYPE\_WOULD\_HAVE\_NO\_PARENT\_TYPE (*parents*)

If *sds* has OWNER granted or denied:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION  
(*new\_type\_in\_sds*)

PARENT\_BASIC\_TYPES\_ARE\_MULTIPLE (*parents*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)

SDS\_IS\_PREDEFINED (*sds*)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, element of *parents*)

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*sds*, *local\_name*)

TYPE\_NAME\_IS\_INVALID (*local\_name*)

### 10.2.12 SDS\_CREATE\_RELATIONSHIP\_TYPE

```

SDS_CREATE_RELATIONSHIP_TYPE (
    sds : Sds_designator,
    forward_local_name : [ Name ],
    forward_category : Category,
    forward_lower_bound : Natural,
    forward_upper_bound : [ Natural ],
    forward_exclusiveness : Exclusiveness,
    forward_stability : Stability,
    forward_duplication : Duplication,
    forward_key_types : [ Key_types_in_sds ],
    reverse_local_name : [ Name ],
    reverse_category : Category,
    reverse_lower_bound : Natural,
    reverse_upper_bound : [ Natural ],
    reverse_exclusiveness : Exclusiveness,
    reverse_stability : Stability,
    reverse_duplication : Duplication,
    reverse_key_types : [ Key_types_in_sds ]
)
    new_forward_type : Link_type_nominator_in_sds,
    new_reverse_type : Link_type_nominator_in_sds

```

SDS\_CREATE\_RELATIONSHIP\_TYPE creates two new non-designation link types *new\_forward\_type* and *new\_reverse\_type*, and their associated types in SDS *new\_forward\_type\_in\_sds* and *new\_reverse\_type\_in\_sds* in the SDS *sds*. The new link types are the reverse of each other.

The operation creates "definition" links from *sds* to *new\_forward\_type\_in\_sds* and *new\_reverse\_type\_in\_sds*, and their reverse "type\_in\_sds" links; the keys of the "definition" links are the system-assigned type identifiers of *new\_forward\_type* and *new\_reverse\_type*. The operation also creates "of\_type" links from *new\_forward\_type\_in\_sds* to *new\_forward\_type* and from *new\_reverse\_type\_in\_sds* to *new\_reverse\_type*, and their reverse "has\_type\_in\_sds" links.

If *forward\_local\_name* is supplied, a "named\_definition" link is created from *sds* to *new\_forward\_type\_in\_sds* with *forward\_local\_name* as key. If *reverse\_local\_name* is supplied, a "named\_definition" link is created from *sds* to *new\_reverse\_type\_in\_sds* with *reverse\_local\_name* as key, together with its reverse "named\_in\_sds" link.

The type identifiers of *new\_forward\_type* and *new\_reverse\_type* are set to implementation-dependent values which identify the types within the PCTE installation.

Two "reverse" links are created between *new\_forward\_type* and *new\_reverse\_type*, one in each direction.

The three definition mode attributes of *new\_forward\_type\_in\_sds* and *new\_reverse\_type\_in\_sds* are set to 19, representing CREATE\_MODE, DELETE\_MODE, and NAVIGATE\_MODE, and their creation or importation times are set to the system time. If *forward\_local\_name* is supplied, the annotation of *new\_forward\_type\_in\_sds* is set to the complete name of the created type; otherwise it is set to the empty string. If *reverse\_local\_name* is supplied, the annotation of *new\_reverse\_type\_in\_sds* is set to the complete name of the created type; otherwise it is set to the empty string.

The category, lower and upper bounds, exclusiveness, stability, and duplication of *new\_forward\_type* are set from *forward\_category*, *forward\_lower\_bound*, *forward\_upper\_bound*, *forward\_exclusiveness*, *forward\_stability*, and *forward\_duplication*, respectively; the category, lower and upper bounds, exclusiveness, stability, and duplication of *new\_reverse\_type* are set from *reverse\_category*, *reverse\_lower\_bound*, *reverse\_upper\_bound*, *reverse\_exclusiveness*, *reverse\_stability*, and *reverse\_duplication*, respectively. Furthermore, if for either link type the cardinality is MANY, for each attribute type of *forward\_key\_types* or *reverse\_key\_types*, a "key\_attribute" link is created from *new\_forward\_type* or *new\_reverse\_type* respectively to that attribute type. The keys of these links correspond to the order of the key attribute types in *forward\_key\_types* or *reverse\_key\_types* starting at 0 and incremented by 1.

The new objects reside on the same volume as *sds*. Their access control lists are built using the default access control list of the calling process, and their confidentiality labels and integrity labels are set to be equal to the current confidentiality context and integrity context, respectively, of the calling process.

For each created object, an "object\_on\_volume" link is created from the volume on which the object resides to the object. The key of the link is the exact identifier of the object.

Write locks of the default mode are obtained on the created objects and links (except the new "object\_on\_volume" links).

## Errors

ACCESS\_ERRORS (*sds*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (element of *forward\_key\_types*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

ACCESS\_ERRORS (element of *reverse\_key\_types*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

KEY\_TYPE\_IS\_BAD (element of *forward\_key\_types*)

KEY\_TYPE\_IS\_BAD (element of *reverse\_key\_types*)

KEY\_TYPES\_ARE\_MULTIPLE (*forward\_key\_types*)

KEY\_TYPES\_ARE\_MULTIPLE (*reverse\_key\_types*)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_DEFINITION\_NAME\_SIZE)

LINK\_TYPE\_CATEGORY\_IS\_BAD (*forward\_link\_type*, (COMPOSITION, EXISTENCE, REFERENCE, IMPLICIT))

LINK\_TYPE\_CATEGORY\_IS\_BAD (*reverse\_link\_type*, (COMPOSITION, EXISTENCE, REFERENCE, IMPLICIT))

LINK\_TYPE\_PROPERTIES\_ARE\_INCONSISTENT (*forward\_category*, *forward\_lower\_bound*, *forward\_upper\_bound*, *forward\_exclusiveness*, *forward\_stability*, *forward\_duplication*)

LINK\_TYPE\_PROPERTIES\_ARE\_INCONSISTENT (*reverse\_category*, *reverse\_lower\_bound*, *reverse\_upper\_bound*, *reverse\_exclusiveness*, *reverse\_stability*, *reverse\_duplication*)

LINK\_TYPE\_PROPERTIES\_AND\_KEY\_TYPES\_ARE\_INCONSISTENT (*forward\_category*, *forward\_lower\_bound*, *forward\_upper\_bound*, *forward\_exclusiveness*, *forward\_stability*, *forward\_duplication*, *forward\_key\_types*)

LINK\_TYPE\_PROPERTIES\_AND\_KEY\_TYPES\_ARE\_INCONSISTENT (*reverse\_category*, *reverse\_lower\_bound*, *reverse\_upper\_bound*, *reverse\_exclusiveness*, *reverse\_stability*, *reverse\_duplication*, *reverse\_key\_types*)

If *sds* has OWNER granted or denied:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION (type in *sds* of *forward\_link\_type*)

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION (type in *sds* of *reverse\_link\_type*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

RELATIONSHIP\_TYPE\_PROPERTIES\_ARE\_INCONSISTENT (*forward\_category*, *forward\_lower\_bound*, *forward\_upper\_bound*, *forward\_exclusiveness*, *forward\_stability*, *forward\_duplication*, *reverse\_category*, *reverse\_lower\_bound*, *reverse\_upper\_bound*, *reverse\_exclusiveness*, *reverse\_stability*, *reverse\_duplication*)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)

SDS\_IS\_PREDEFINED (*sds*)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, element of *forward\_key\_types*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, element of *reverse\_key\_types*)

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*forward\_local\_name*)

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*reverse\_local\_name*)

TYPE\_NAME\_IS\_INVALID (*forward\_local\_name*)

TYPE\_NAME\_IS\_INVALID (*reverse\_local\_name*)

### 10.2.13 SDS\_CREATE\_STRING\_ATTRIBUTE\_TYPE

```
SDS_CREATE_STRING_ATTRIBUTE_TYPE (
    sds           : Sds_designator,
    local_name    : [ Name ],
    initial_value : [ String ],
    duplication   : Duplication
)
new_type       : Attribute_type_nominator_in_sds
```

SDS\_CREATE\_STRING\_ATTRIBUTE\_TYPE creates a new string attribute type *new\_type*, and its associated attribute type in SDS *new\_type\_in\_sds* in the SDS *sds*.

The operation creates a "definition" link from *sds* to *new\_type\_in\_sds*; the key of the link is the system-assigned type identifier of *new\_type*. The operation also creates an "of\_type" link from *new\_type\_in\_sds* to *new\_type*.

If *local\_name* is supplied, a "named\_definition" link is created from *sds* to *new\_type\_in\_sds* with *local\_name* as key, together with its reverse "named\_in\_sds" link.

The type identifier of *new\_type* is set to an implementation-defined value which identifies the type within the PCTE installation. The duplication of *new\_type* is set to *duplication*. The string initial value of *new\_type* is set to *initial\_value* if supplied, and otherwise to the empty string.

The three definition mode attributes of *new\_type\_in\_sds* are set to 12, representing READ\_MODE and WRITE\_MODE, and its creation or importation time is set to the system

time. If *local\_name* is supplied, the annotation of *new\_type\_in\_sds* is set to the complete name of the created type; otherwise it is set to the empty string.

The new objects reside on the same volume as *sds*. Their access control lists are built using the default atomic ACL and the default object owner of the calling process, and their confidentiality labels and integrity labels are set to be equal to the current confidentiality context and integrity context, respectively, of the calling process.

For each created object, an "object\_on\_volume" link is created from the volume on which the object resides to the object. The key of the link is the exact identifier of the object.

Write locks of the default mode are obtained on the created objects and links except the new "object\_on\_volume" links.

### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, MODIFY, APPEND\_LINKS)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_DEFINITION\_NAME\_SIZE)

If *sds* has OWNER granted or denied:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION  
(*new\_type\_in\_sds*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)

SDS\_IS\_PREDEFINED (*sds*)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*sds*, *local\_name*)

TYPE\_NAME\_IS\_INVALID (*local\_name*)

VALUE\_LIMIT\_ERRORS (*initial\_value*)

### 10.2.14 SDS\_CREATE\_TIME\_ATTRIBUTE\_TYPE

```
SDS_CREATE_TIME_ATTRIBUTE_TYPE (
    sds           : Sds_designator,
    local_name    : [ Name ],
    initial_value : [ Time ],
    duplication   : Duplication
)
new_type       : Attribute_type_nominator_in_sds
```

SDS\_CREATE\_TIME\_ATTRIBUTE\_TYPE creates a new time attribute type *new\_type*, and its associated attribute type in SDS *new\_type\_in\_sds* in the SDS *sds*.

The operation creates a "definition" link from *sds* to *new\_type\_in\_sds*; the key of the link is the system-assigned type identifier of *new\_type*. The operation also creates an "of\_type" link from *new\_type\_in\_sds* to *new\_type*.

If *local\_name* is supplied, a "named\_definition" link is created from *sds* to *new\_type\_in\_sds* with *local\_name* as key, together with its reverse "named\_in\_sds" link.

The type identifier of *new\_type* is set to an implementation-defined value which identifies the type within the PCTE installation. The duplication of *new\_type* is set to *duplication*. The time

initial value of *new\_type* is set to *initial\_value* if supplied, and otherwise to 1980-01-01T00:00:00Z.

The three definition mode attributes of *new\_type\_in\_sds* are set to 12, representing READ\_MODE and WRITE\_MODE, and its creation or importation time is set to the system time. If *local\_name* is supplied, the annotation of *new\_type\_in\_sds* is set to the complete name of the created type; otherwise it is set to the empty string.

The new objects reside on the same volume as *sds*. Their access control lists are built using the default atomic ACL and the default object owner of the calling process, and their confidentiality labels and integrity labels are set to be equal to the current confidentiality context and integrity context, respectively, of the calling process.

For each created object, an "object\_on\_volume" link is created from the volume on which the object resides to the object. The key of the link is the exact identifier of the object.

Write locks of the default mode are obtained on the created objects and links except the new "object\_on\_volume" links.

### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, MODIFY, APPEND\_LINKS)

ATTRIBUTE\_VALUE\_LIMIT\_WOULD\_BE\_EXCEEDED (*initial\_value*)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_DEFINITION\_NAME\_SIZE)

If *sds* has OWNER granted or denied:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION  
(*new\_type\_in\_sds*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)

SDS\_IS\_PREDEFINED (*sds*)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*sds*, *local\_name*)

TYPE\_NAME\_IS\_INVALID (*local\_name*)

VALUE\_LIMIT\_ERRORS (*initial\_value*)

### 10.2.15 SDS\_GET\_NAME

```
SDS_GET_NAME (
    sds      : Sds_designator
)
    name    : Name
```

SDS\_GET\_NAME returns the name of the SDS *sds*.

The returned name *name* is the key of the "known\_sds" link from the SDS directory to *sds*.

A read lock of the default mode is obtained on that link.

### Errors

ACCESS\_ERRORS (the SDS directory, ATOMIC, READ, READ\_LINKS)

SDS\_IS\_PREDEFINED (*sds*)



SDS\_IS\_UNKNOWN (*sds*)

### 10.2.16 SDS\_IMPORT\_ATTRIBUTE\_TYPE

```
SDS_IMPORT_ATTRIBUTE_TYPE (
    to_sds      : Sds_designator,
    from_sds    : Sds_designator,
    type        : Attribute_type_nominator_in_sds,
    local_name   : [ Name ]
)
```

SDS\_IMPORT\_ATTRIBUTE\_TYPE imports the attribute type *type* from the SDS *from\_sds* to the SDS *to\_sds*, along with the associated enumerational types if *type* is an enumeration attribute type.

If *type* is an enumeration attribute type, all its enumerational types are implicitly imported, if not already in *to\_sds*. The implicitly imported enumerational types have the same images as in *from\_sds*, but do not have local names in *to\_sds*.

The operation creates an attribute type in SDS *type\_in\_sds* in *to\_sds* associated with *type*. If *type* is an enumeration attribute type, it also creates an enumerational type in SDS in *sds* associated with each enumerational type of *type* (unless one already exists). For each of the created types in SDS a "definition" link is created from *to\_sds* whose key is the type identifier of the associated type.

An "of\_type" link from each new type in SDS to its associated type and its reverse "has\_type\_in\_sds" link are created.

If *local\_name* is supplied, or if *type* has a local name in *from\_sds*, a "named\_definition" link from *to\_sds* to *type\_in\_sds* and its reverse "named\_in\_sds" link are created. The key of the "named\_definition" link is *local\_name* if supplied, otherwise the local name of *type* in *from\_sds*.

Each of the three definition mode attributes of *type\_in\_sds* is set to the export mode of the corresponding type in SDS in *from\_sds*.

The creation or importation time of each new type in SDS is set to the system time.

The annotation of each new type in SDS is the same as the annotation of the corresponding type in SDS in *from\_sds*.

The new types in SDS reside on the same volume as *to\_sds*. Their access control lists are built using the default atomic ACL and default object owner of the calling process, and their confidentiality labels and integrity labels are set to be equal to the current confidentiality context and integrity context, respectively, of the calling process.

An "object\_on\_volume" link is created to each new type in SDS from the volume on which it resides. The key of the link is the exact identifier of the new type in SDS.

Read locks of the default mode are obtained on the types in SDS in *from\_sds*. Write locks of the default mode are obtained on the new types in SDS and links (except the new "object\_on\_volume" links).

#### Errors

ACCESS\_ERRORS (*from\_sds*, ATOMIC, READ, NAVIGATE)

ACCESS\_ERRORS (*to\_sds*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (the type in SDS associated with *type* in *from\_sds*, ATOMIC, READ, EXPLOIT\_SCHEMA)

ACCESS\_ERRORS (the type in SDS associated with *type* in *from\_sds*, ATOMIC, READ, READ\_ATTRIBUTES)

ACCESS\_ERRORS (the "type" object associated with *type*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

If *type* is an enumeration attribute type, for each enumerational type in SDS *S* associated with *type* in *from\_sds*:

ACCESS\_ERRORS (*S*, ATOMIC, READ, EXPLOIT\_SCHEMA)

If *type* is an enumeration attribute type, for each enumerational type *E* associated with *type* not already present in *to\_sds*:

ACCESS\_ERRORS (*E*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

If *type* is an enumeration attribute type:

IMAGE\_IS\_DUPLICATED (enumerational types of *type*, *to\_sds*)

If *to\_sds* has OWNER granted or denied:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION (*type\_in\_sds*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*to\_sds*)

SDS\_IS\_PREDEFINED (*sds*)

SDS\_IS\_UNKNOWN (*to\_sds*)

SDS\_IS\_UNKNOWN (*from\_sds*)

TYPE\_IS\_ALREADY\_KNOWN\_IN\_SDS (*type*, *to\_sds*)

If *local\_name* is supplied:

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*to\_sds*, *local\_name*)

If *local\_name* is not supplied:

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*to\_sds*, local name of *type* in *from\_sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*from\_sds*, *type*)

TYPE\_NAME\_IS\_INVALID (*local\_name*)

### 10.2.17 SDS\_IMPORT\_ENUMERAL\_TYPE

```
SDS_IMPORT_ENUMERAL_TYPE (
    to_sds      : Sds_designator,
    from_sds    : Sds_designator,
    type        : Enumeral_type_nominator_in_sds,
    local_name  : [ Name ]
)
```

SDS\_IMPORT\_ENUMERAL\_TYPE imports the enumerational type *type* from the SDS *from\_sds* to the SDS *to\_sds*.

The operation creates an enumerational type in SDS *type\_in\_sds* in *to\_sds* associated with *type*. A "definition" link is created from *to\_sds* to *type\_in\_sds* whose key is the type identifier of *type*, together with its reverse "in\_sds" link.

An "of\_type" link is created from *type\_in\_sds* to *type*, together with its reverse "has\_type\_in\_sds" link.

If *local\_name* is supplied, or if *type* has a local name in *from\_sds*, a "named\_definition" link is created from *to\_sds* to *type\_in\_sds*, together with its reverse "named\_in\_sds" link. The key of the "named\_definition" link is *local\_name* if supplied, otherwise the local name of *type* in *from\_sds*.

The creation or importation time of *type\_in\_sds* is set to the system time.

The annotation of *type\_in\_sds* is the same as the annotation of the corresponding type in SDS in *from\_sds*.

*type\_in\_sds* resides on the same volume as *to\_sds*. Its access control lists are built using the default atomic ACL and the default object owner of the calling process, and its confidentiality label and integrity label is set to be equal to the current confidentiality context and integrity context, respectively, of the calling process.

An "object\_on\_volume" link is created to *type\_in\_sds* from the volume on which it resides. The key of the link is the exact identifier of *type\_in\_sds*.

Read locks of the default mode are obtained on the type in SDS in *from\_sds*. Write locks of the default mode are obtained on *type\_in\_sds* and the created links (except the new "object\_on\_volume" link).

## Errors

ACCESS\_ERRORS (*from\_sds*, ATOMIC, READ, NAVIGATE)

ACCESS\_ERRORS (*to\_sds*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (the type in SDS associated with *type* in *from\_sds*, ATOMIC, READ, EXPLOIT\_SCHEMA)

ACCESS\_ERRORS (the type in SDS associated with *type* in *from\_sds*, ATOMIC, READ, READ\_ATTRIBUTES)

ACCESS\_ERRORS (the "type" object associated with *type*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

If *to\_sds* has OWNER granted or denied:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION (*type\_in\_sds*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*to\_sds*)

SDS\_IS\_PREDEFINED (*to\_sds*)

SDS\_IS\_UNKNOWN (*to\_sds*)

SDS\_IS\_UNKNOWN (*from\_sds*)

TYPE\_IS\_ALREADY\_KNOWN\_IN\_SDS (*type*, *to\_sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*from\_sds*, *type*)

If *local\_name* is supplied:

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*to\_sds*, *local\_name*)

If *local\_name* is not supplied:

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*to\_sds*, local name of *type* in *from\_sds*)

TYPE\_NAME\_IS\_INVALID (*local\_name*)

### 10.2.18 SDS\_IMPORT\_LINK\_TYPE

```

SDS_IMPORT_LINK_TYPE (
    to_sds      : Sds_designator,
    from_sds    : Sds_designator,
    type        : Link_type_nominator_in_sds,
    local_name  : [ Name ]
)

```

SDS\_IMPORT\_LINK\_TYPE imports the link type *type* from the SDS *from\_sds* to the SDS *to\_sds*.

All the key attribute types of *type*, and its reverse link type with all its key attributes, are implicitly imported, if not already in *to\_sds*. The imported link types have the same key attributes as in *from\_sds*. The link and the key attribute types implicitly imported do not have local names assigned to them within *to\_sds*.

The importation of a type (either explicitly or implicitly) results in the creation of a type in SDS in *to\_sds* associated with the imported type, with a "definition" link from *from\_sds* whose key is the type identifier of the imported type. An "of\_type" link from the new type in SDS to the imported type and its reverse "has\_type\_in\_sds" link are created.

If *local\_name* is supplied or if *type* has a name in *from\_sds*, a "named\_definition" link is created from *to\_sds* to the new link type in SDS associated with *type*, together with its reverse "named\_in\_sds" link. The key of the "named\_definition" link is *local\_name* if supplied, otherwise it is the local name of *type* in the SDS *from\_sds*.

Each of the three definition mode attributes of each new type in SDS is set to the export mode of the corresponding type in SDS in *from\_sds*.

The creation or importation time of each new type in SDS is set to the system time.

The annotation of each new type in SDS is the same as the annotation of the corresponding type in SDS in *from\_sds*.

The new types in SDS reside on the same volume as *to\_sds*. Their access control lists are built using the default atomic ACL and default object owner of the calling process, and their confidentiality labels and integrity labels are set to be equal to the current confidentiality context and integrity context, respectively, of the calling process.

An "object\_on\_volume" link is created to each new type in SDS from the volume on which it resides. The key of the link is the exact identifier of the new type in SDS.

Read locks of the default mode are obtained on the types in SDS in *from\_sds*. Write locks of the default mode are obtained on the new types in SDS and links (except the new "object\_on\_volume" links).

#### Errors

ACCESS\_ERRORS (*from\_sds*, ATOMIC, READ, NAVIGATE)

ACCESS\_ERRORS (*to\_sds*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (type in SDS associated with *type* in *from\_sds*, ATOMIC, READ, EXPLOIT\_SCHEMA)

ACCESS\_ERRORS (the type in SDS associated with *type* in *from\_sds*, ATOMIC, READ, READ\_ATTRIBUTES)

For each attribute type in SDS A associated with a key attribute type of *type*:

ACCESS\_ERRORS (A, ATOMIC, READ, EXPLOIT\_SCHEMA)

ACCESS\_ERRORS (the "type" object associated with *type*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

For each key attribute type K of *type* not already present in *to\_sds*:

ACCESS\_ERRORS (K, ATOMIC, CHANGE, APPEND\_IMPLICIT)

If *type* has a reverse link type R not already present in *to\_sds*:

ACCESS\_ERRORS (R, ATOMIC, CHANGE, APPEND\_IMPLICIT)

For each key attribute type K1 of R not already present in *to\_sds*:

ACCESS\_ERRORS (K1, ATOMIC, CHANGE, APPEND\_IMPLICIT)

If *to\_sds* has OWNER granted or denied:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION (*type\_in\_sds*)

If *to\_sds* has OWNER granted or denied and *link\_type* has a reverse link type:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION (reverse of *type\_in\_sds*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*to\_sds*)

SDS\_IS\_PREDEFINED (*to\_sds*)

SDS\_IS\_UNKNOWN (*to\_sds*)

SDS\_IS\_UNKNOWN (*from\_sds*)

TYPE\_IS\_ALREADY\_KNOWN\_IN\_SDS (*type*, *to\_sds*)

If *local\_name* is supplied:

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*to\_sds*, *local\_name*)

If *local\_name* is not supplied:

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*to\_sds*, local name of *type* in *from\_sds*)

TYPE\_NAME\_IS\_INVALID (*local\_name*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*from\_sds*, *type*)

#### 10.2.19 SDS\_IMPORT\_OBJECT\_TYPE

```
SDS_IMPORT_OBJECT_TYPE (
    to_sds      : Sds_designator,
    from_sds    : Sds_designator,
    type        : Object_type_nominator_in_sds,
    local_name  : [ Name ]
)
```

SDS\_IMPORT\_OBJECT\_TYPE imports the object type *type* from the SDS *from\_sds* to the SDS *to\_sds*.

The importation of an object type implies the implicit importation of all its ancestor types if not already in *to\_sds*. The attribute and link types applied to the explicitly or implicitly imported types are not imported, nor is the notion of their application. The object types implicitly imported do not have a local name assigned to them within *to\_sds*.

The importation of an object type (either explicit or implicit) results in the creation of an object type in SDS in *to\_sds* with a "definition" link from *to\_sds* whose key is the type identifier of the

imported type. An "of\_type" link from the new object type in SDS to the imported type and its reverse "has\_type\_in\_sds" link are created.

If *local\_name* is supplied or if the imported *type* has a name in the originating SDS, a "named\_definition" link is created from *to\_sds* to the new object type in SDS associated with *link\_type*, together with its reverse "named\_in\_sds" link. The key of the "named\_definition" link is *local\_name* if supplied, otherwise it is the local name of *type* in *from\_sds*.

Each of the three definition mode attributes of each new type in SDS is set to the export mode of the corresponding type in SDS in *from\_sds*.

The creation or importation time of each new type in SDS is set to the system time.

The annotation of each new type in SDS is the same as the annotation of the corresponding type in SDS in *from\_sds*.

The new types in SDS reside on the same volume as *to\_sds*. Their access control lists are built using the default atomic ACL and default object owner of the calling process, and their confidentiality labels and integrity labels are set to be equal to the current confidentiality context and integrity context, respectively, of the calling process.

An "object\_on\_volume" link is created to each new type in SDS from the volume on which it resides. The key of the link is the exact identifier of the new type in SDS.

Read locks of the default mode are obtained on the types in SDS in *from\_sds*. Write locks of the default mode are obtained on the new types in SDS and links (except the new "object\_on\_volume" links).

## Errors

ACCESS\_ERRORS (*from\_sds*, ATOMIC, READ, NAVIGATE)

ACCESS\_ERRORS (*to\_sds*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (the type in SDS associated with /type/ in /from\_sds/, ATOMIC, READ, READ\_ATTRIBUTES)

For each object type in SDS S associated with the ancestor types of *type*:

ACCESS\_ERRORS (S, ATOMIC, READ, EXPLOIT\_SCHEMA)

ACCESS\_ERRORS (object type in SDS associated with *type* in *from\_sds*, ATOMIC, READ, EXPLOIT\_SCHEMA)

ACCESS\_ERRORS (the "type" object associated with *type*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

For each ancestor object type A of *type* not already present in *to\_sds*:

ACCESS\_ERRORS(A, ATOMIC, CHANGE, APPEND\_IMPLICIT)

If *to\_sds* has OWNER granted or denied:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION (*type\_in\_sds*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*to\_sds*)

SDS\_IS\_PREDEFINED (*to\_sds*)

SDS\_IS\_UNKNOWN (*to\_sds*)

SDS\_IS\_UNKNOWN (*from\_sds*)

TYPE\_IS\_ALREADY\_KNOWN\_IN\_SDS (*type*, *to\_sds*)

If *local\_name* is supplied:

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*to\_sds*, *local\_name*)

If *local\_name* is not supplied:

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*to\_sds*, local name of *type* in *from\_sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*from\_sds*, *type*)

TYPE\_NAME\_IS\_INVALID (*local\_name*)

### 10.2.20 SDS\_INITIALIZE

```
SDS_INITIALIZE (
    sds      : Sds_designator,
    name     : Name
)
```

SDS\_INITIALIZE establishes the SDS *sds* as a known SDS by creating a "known\_sds" link with key *name* from the master of the SDS directory to *sds*.

A read lock of the default mode is obtained on *sds* and a write lock of the default mode is obtained on the created link.

#### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

ACCESS\_ERRORS (master of the SDS directory, ATOMIC, MODIFY, APPEND\_LINKS)

If *sds* has successors or predecessors:

ACCESS\_ERRORS (successor or predecessor of *sds*, ATOMIC, SYSTEM\_ACCESS)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_DEFINITION\_NAME\_SIZE)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_KNOWN (*sds*)

SDS\_IS\_NOT\_EMPTY\_NOR\_VERSION (*sds*)

SDS\_IS\_PREDEFINED (*sds*)

SDS\_NAME\_IS\_DUPLICATE (*name*)

SDS\_NAME\_IS\_INVALID (*name*)

### 10.2.21 SDS\_REMOVE

```
SDS_REMOVE (
    sds      : Sds_designator
)
```

SDS\_REMOVE removes the SDS *sds* from the set of known SDSs.

The "known\_sds" link to *sds* from the SDS directory is deleted. If that link is the last composition or existence link to *sds*, then the "*sds*" object *sds* is deleted. In that case, the "object\_on\_volume" link from the volume on which *sds* was residing to *sds* is also deleted.

A read lock of the default mode is obtained on *sds* if it is not deleted; a write lock otherwise. Write locks of the default mode are obtained on the deleted links except the deleted "object\_on\_volume" link.



**Errors**

ACCESS\_ERRORS (the SDS directory, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (*sds*, ATOMIC, CHANGE, WRITE\_IMPLICIT)

If the conditions hold for deletion of the "sds" object *sds*:

ACCESS\_ERRORS (*sds*, COMPOSITE, MODIFY, DELETE)

If *sds* has predecessors or successors:

ACCESS\_ERRORS (predecessor or successor of *sds*, ATOMIC, SYSTEM\_ACCESS)

OBJECT\_HAS\_LINKS\_PREVENTING\_DELETION (*sds*)

OBJECT\_IS\_IN\_USE\_FOR\_DELETE (*sds*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)

SDS\_IS\_NOT\_EMPTY\_NOR\_VERSION (*sds*)

SDS\_IS\_PREDEFINED (*sds*)

SDS\_IS\_UNKNOWN (*sds*)

**10.2.22 SDS\_REMOVE\_DESTINATION**

```
SDS_REMOVE_DESTINATION(
    sds           : Sds_designator,
    link_type      : Link_type_nominator_in_sds,
    object_type    : Object_type_nominator_in_sds
)
```

SDS\_REMOVE\_DESTINATION removes the object type in SDS *object\_type\_in\_sds* associated with *object\_type* in the SDS *sds* from the destination object types of the link type in SDS *link\_type\_in\_sds* associated with *link\_type* in *sds*.

As a result, the "in\_destination\_set" link from *link\_type\_in\_sds* to *object\_type\_in\_sds* and its reverse "is\_destination\_of" link are deleted.

If *link\_type* has a reverse link type *reverse*, then *reverse* is unapplied (see SDS\_UNAPPLY\_LINK\_TYPE) and the "in\_link\_set" link existing between *object\_type\_in\_sds* and the "link\_type\_in\_sds" object *reverse\_link\_type\_in\_sds* associated with *reverse* in *sds* is deleted.

Write locks of the default mode are obtained on the deleted links.

**Errors**

ACCESS\_ERRORS (*sds*, ATOMIC, READ, NAVIGATE)

ACCESS\_ERRORS (*object\_type\_in\_sds*, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (*link\_type\_in\_sds*, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (*reverse\_link\_type\_in\_sds*, ATOMIC, MODIFY, WRITE\_LINKS)

OBJECT\_TYPE\_IS\_NOT\_IN\_DESTINATION\_SET (*link\_type*, *object\_type*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)

SDS\_IS\_PREDEFINED (*sds*)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *link\_type*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *object\_type*)

### 10.2.23 SDS\_REMOVE\_TYPE

```
SDS_REMOVE_TYPE (
    sds      : Sds_designator,
    type     : Type_nominator_in_sds
)
```

SDS\_REMOVE\_TYPE removes from the SDS *sds* the type in SDS *type\_in\_sds* associated with *type* in *sds*.

When a link type is removed from an SDS, the type in SDS *reverse\_link\_type\_in\_sds* associated with the reverse type *reverse\_type* of *type* (if any) is also removed from that SDS.

If *type\_in\_sds* is the last type in SDS associated with *type* in any SDS, then *type* is also removed. If *reverse\_link\_type\_in\_sds* exists and is the last type in SDS associated with *reverse\_type* in any SDS, then *reverse\_type* is also removed. The type identifier of a removed type is never reassigned.

A type can be removed while there are instances of the type in the object base. Instances of the removed type are not directly affected by this operation: objects, attributes and links retain the type properties of the type. The description of the type is however lost; the implications are:

- no instances of a removed type can be created;
- attributes of a removed type are inaccessible. They can however still be consulted or reset to their initial value using the system-generated type identifier of the deleted type (see 23.1.2.5);
- links of a removed type can still be navigated through or deleted using the type identifier of the deleted type;
- objects of a removed type can still be accessed as instances of visible ancestor types of the removed type.

The *removal* of a type in SDS consists of the deletion of the "definition" link and "named\_definition" link (if any) between the SDS and the type in SDS. The deletion of the "definition" link may result in the deletion of the type in SDS.

The deletion of a type in SDS also entails the deletion of the "of\_type" link from the type in SDS. In the case where this link is the last "of\_type" link to the type (i.e. if the last occurrence of the type in SDS is to be deleted) the type is also removed.

In turn, the removal of a type entails the loss of all the typing information held on its associated type object (such as the "parent\_type", "key\_attribute", "reverse" or "enumerals" links starting from that object) and may imply the deletion of the type itself (if the "of\_type" link to be deleted is the last existence link to it and if there are no composition links to it).

For each deleted object, the "object\_on\_volume" link from the volume on which the deleted object was residing to the deleted object is also deleted. A write lock of the default mode is obtained on *sds* and on the deleted objects and links (except the deleted "object\_on\_volume" links).

**Errors**

ACCESS\_ERRORS (*sds*, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (*type\_in\_sds*, ATOMIC, MODIFY, (WRITE\_LINKS, WRITE\_IMPLICIT))

If the deletion of a type or type in SDS is implied:

ACCESS\_ERRORS (destination object of a link from the type or type in SDS to be deleted which has an implicit reverse link, ATOMIC, CHANGE, WRITE\_IMPLICIT)

If conditions hold for the deletion of *type\_in\_sds*:

ACCESS\_ERRORS (*type\_in\_sds*, COMPOSITE, MODIFY, DELETE)

ACCESS\_ERRORS ("type" object associated with *type*, ATOMIC, CHANGE, WRITE\_IMPLICIT)

ACCESS\_ERRORS (*type\_in\_sds*, COMPOSITE, MODIFY, DELETE)

If the conditions for the deletion of the "type" object T associated with *type* are satisfied:

ACCESS\_ERRORS (T, COMPOSITE, MODIFY, DELETE)

If T is an object type, for each parent type P of T:

ACCESS\_ERRORS (P, ATOMIC, CHANGE, WRITE\_IMPLICIT)

If T is a link type, for each key attribute type K of T:

ACCESS\_ERRORS (K, ATOMIC, CHANGE, WRITE\_IMPLICIT)

If T is a link type with a reverse link type R:

ACCESS\_ERRORS (R, COMPOSITE, MODIFY, DELETE)

For each key attribute type K1 of R:

ACCESS\_ERRORS (K1, ATOMIC, CHANGE, WRITE\_IMPLICIT)

If T is an enumeration attribute type, for each associated enumerational type E:

ACCESS\_ERRORS (E, ATOMIC, CHANGE, WRITE\_IMPLICIT)

If the deletion of a type or type in SDS is implied:

OBJECT\_HAS\_LINKS\_PREVENTING\_DELETION (type or type in SDS to be deleted)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)

SDS\_IS\_PREDEFINED (*sds*)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_HAS\_DEPENDENCIES (*sds*, *type*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *type*)

**10.2.24 SDS\_SET\_ENUMERAL\_TYPE\_IMAGE**

```
SDS_SET_ENUMERAL_TYPE_IMAGE (
    sds      : Sds_designator,
    type     : Enumeral_type_nominator_in_sds,
    image    : [ Text ]
)
```

SDS\_SET\_ENUMERAL\_TYPE\_IMAGE sets the image of the enumerational type in SDS *type\_in\_sds* associated with the enumerational type *type* in the SDS *sds* to the text value *image*, if supplied; if *image* is not supplied, the image of *type\_in\_sds* is set to the empty text value.

A write lock of the default mode is obtained on *type\_in\_sds*.

### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, READ, NAVIGATE)

ACCESS\_ERRORS (*type\_in\_sds*, ATOMIC, MODIFY, WRITE\_ATTRIBUTES)

IMAGE\_IS\_ALREADY\_ASSOCIATED (*image*, *sds*, *type*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)

SDS\_IS\_PREDEFINED (*sds*)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *type*)

## 10.2.25 SDS\_SET\_TYPE\_MODES

```
SDS_SET_TYPE_MODES (
    sds           : Sds_designator,
    type          : Object_type_nominator_in_sds | Attribute_type_nominator_in_sds |
                  Link_type_nominator_in_sds,
    usage_mode    : [ Definition_mode_values ],
    export_mode   : [ Definition_mode_values ]
)
```

SDS\_SET\_TYPE\_MODES sets the usage mode and export mode of the type in SDS *type\_in\_sds* associated with the type *type* in the SDS *sds* to the values of *usage\_mode* and *export\_mode* respectively, if supplied; if, for either parameter, no value is supplied, then the corresponding value is left unchanged.

If *type* is a link type with a reverse link type *reverse*, the usage mode and export mode of the link type in *sds reverse\_type\_in\_sds* associated with *reverse* in *sds* are set (or not) in the same way.

Write locks of the default mode are obtained on the modified "type\_in\_sds" objects.

### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, READ, NAVIGATE)

ACCESS\_ERRORS (*type\_in\_sds*, COMPOSITE, MODIFY, WRITE\_ATTRIBUTES)

If *type* has a reverse link type in SDS *reverse\_type\_in\_sds*:

ACCESS\_ERRORS (*reverse\_type\_in\_sds*, ATOMIC, MODIFY, WRITE\_ATTRIBUTES)

DEFINITION\_MODE\_VALUE\_WOULD\_BE\_INVALID (*export\_mode*, *type*)

DEFINITION\_MODE\_VALUE\_WOULD\_BE\_INVALID (*usage\_mode*, *type*)

MAXIMUM\_USAGE\_MODE\_WOULD\_BE\_EXCEEDED (*type*, *usage\_mode*)

MAXIMUM\_USAGE\_MODE\_WOULD\_BE\_EXCEEDED (*type*, *export\_mode*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)

SDS\_IS\_PREDEFINED (*sds*)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *type*)

#### 10.2.26 SDS\_SET\_TYPE\_NAME

```
SDS_SET_TYPE_NAME (
    sds           : Sds_designator,
    type          : Type_nominator_in_sds,
    local_name    : [ Name ]
)
```

SDS\_SET\_TYPE\_NAME sets the local name of the type in SDS *type\_in\_sds* associated with the type *type* in the SDS *sds* to *local\_name*, if supplied, and otherwise deletes the local name of *type\_in\_sds* (if any).

If *local\_name* is supplied, a "named\_definition" link is created from *sds* to *type\_in\_sds*. *local\_name* is used as the key of the new link.

If *type\_in\_sds* already had a local name, the corresponding "named\_definition" link is deleted.

Write locks of the default mode are obtained on the deleted links (if any) and write locks of the default mode are obtained on the created links (if any).

#### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, READ, NAVIGATE)

If *type\_in\_sds* already has a local name:

ACCESS\_ERRORS (*sds*, ATOMIC, MODIFY, WRITE\_LINKS)

If *local\_name* is supplied:

ACCESS\_ERRORS (*sds*, ATOMIC, MODIFY, APPEND\_LINKS)

If *type\_in\_sds* already has a local name:

ACCESS\_ERRORS (*type\_in\_sds*, ATOMIC, CHANGE, WRITE\_IMPLICIT)

If *local\_name* is supplied:

ACCESS\_ERRORS (*type\_in\_sds*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)

SDS\_IS\_PREDEFINED (*sds*)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *type*)

If *local\_name* is supplied:

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*sds*, *local\_name*)

TYPE\_NAME\_IS\_INVALID (*local\_name*)

#### 10.2.27 SDS\_UNAPPLY\_ATTRIBUTE\_TYPE

```
SDS_UNAPPLY_ATTRIBUTE_TYPE (
    sds           : Sds_designator,
    attribute_type : Attribute_type_nominator_in_sds,
    type          : Object_type_nominator_in_sds | Link_type_nominator_in_sds
)
```

SDS\_UNAPPLY\_ATTRIBUTE\_TYPE removes the application of the attribute type in SDS *attribute\_type\_in\_sds* associated with the attribute type *attribute\_type* in the SDS *sds* from the type in SDS *type\_in\_sds* associated with the object or link type *type* in *sds*.

The "in\_attribute\_set" link between *type\_in\_sds* and *attribute\_type\_in\_sds* and its reverse "is\_attribute\_of" link are deleted.

Write locks of the default mode are obtained on the deleted links.

#### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, READ, NAVIGATE)

ACCESS\_ERRORS (*type\_in\_sds*, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (*attribute\_type\_in\_sds*, ATOMIC, MODIFY, WRITE\_LINKS)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)

SDS\_IS\_PREDEFINED (*sds*)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_NOT\_APPLIED (*sds*, *attribute\_type*, *type*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *attribute\_type*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *type*)

### 10.2.28 SDS\_UNAPPLY\_LINK\_TYPE

```
SDS_UNAPPLY_LINK_TYPE (
    sds           : Sds_designator,
    link_type      : Link_type_nominator_in_sds,
    object_type    : Object_type_nominator_in_sds
)
```

SDS\_UNAPPLY\_LINK\_TYPE removes the application of the link type in SDS *link\_type\_in\_sds* associated with the link type *link\_type* in the SDS *sds* from the object type in SDS *object\_type\_in\_sds* associated with the object type *object\_type* in *sds*.

The "in\_link\_set" link between *object\_type\_in\_sds* and *link\_type\_in\_sds* and its reverse "is\_link\_of" link are deleted.

If *link\_type* has a reverse link type *reverse*, then *object\_type* is removed from the destination object types of *reverse* (see SDS\_REMOVE\_DESTINATION) and the "in\_destination\_set" link between the link type in SDS *reverse\_link\_type\_in\_sds* associated with *reverse* in *sds* and *object\_type\_in\_sds* is deleted.

Write locks of the default mode are obtained on the deleted links.

#### Errors

ACCESS\_ERRORS (*object\_type\_in\_sds*, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (*link\_type\_in\_sds*, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (*reverse\_link\_type\_in\_sds*, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (*sds*, ATOMIC, READ, NAVIGATE)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SCHEMA\_UPDATE)

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds*)  
 SDS\_IS\_PREDEFINED (*sds*)  
 SDS\_IS\_UNKNOWN (*sds*)  
 TYPE\_IS\_NOT\_APPLIED (*sds*, *link\_type*, *object\_type*)  
 TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *link\_type*)  
 TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *object\_type*)

### 10.3 SDS usage operations

#### 10.3.1 SDS\_GET\_ATTRIBUTE\_TYPE\_PROPERTIES

```

SDS_GET_ATTRIBUTE_TYPE_PROPERTIES (
    sds           : Sds_designator,
    type          : Attribute_type_nominator_in_sds
)
    duplication   : Duplication,
    value_type    : Value_type,
    initial_value : Attribute_value
  
```

SDS\_GET\_ATTRIBUTE\_TYPE\_PROPERTIES returns the duplication, value type identifier, and initial value of the attribute type in SDS identified by *type* in the SDS *sds*.

Read locks of the default mode are obtained on *sds* and on the "type" and "type\_in\_sds" objects associated with *type* in *sds*.

#### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, READ, READ\_LINKS)

ACCESS\_ERRORS ("type" and "type\_in\_sds" objects associated with *type*, ATOMIC, READ, READ\_ATTRIBUTES)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *type*)

#### 10.3.2 SDS\_GET\_ENUMERAL\_TYPE\_IMAGE

```

SDS_GET_ENUMERAL_TYPE_IMAGE (
    sds       : Sds_designator,
    type      : Enumeral_type_nominator_in_sds
)
    image : Text
  
```

SDS\_GET\_ENUMERAL\_TYPE\_IMAGE returns the image *image* of the enumeral type in SDS identified by *type* in the SDS *sds*.

Read locks of the default mode are obtained on *sds* and on the "type" and "type\_in\_sds" objects associated with *type* in *sds*.

#### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, READ, READ\_LINKS)

ACCESS\_ERRORS ("type" and "type\_in\_sds" objects associated with *type*, ATOMIC, READ, READ\_ATTRIBUTES)



SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *type*)

### 10.3.3 SDS\_GET\_ENUMERAL\_TYPE\_POSITION

```
SDS_GET_ENUMERAL_TYPE_POSITION (
    sds      : Sds_designator,
    type1    : Enumeral_type_nominator_in_sds,
    type2    : Attribute_type_nominator_in_sds
)
    position : Natural
```

SDS\_GET\_ENUMERAL\_TYPE\_POSITION returns the position *position* of the enumerai type in SDS identified by *type1* in the SDS *sds*, in the value type of the attribute type in SDS identified by *type2* in *sds*, i.e. the key of the "enumerai" link from the "type" object associated with *type2* to the "type" object associated with *type1*.

Read locks of the default mode are obtained on the "type" and "type\_in\_sds" objects associated with *type1* and *type2* in *sds* and on the "enumerai" link.

#### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, READ, READ\_LINKS)

ACCESS\_ERRORS ("type" and "type\_in\_sds" objects associated with *type1* and *type2*, ATOMIC, READ, READ\_ATTRIBUTES)

ENUMERAL\_TYPE\_IS\_NOT\_IN\_ATTRIBUTE\_VALUE\_TYPE (*type1*, *type2*)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *type1*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *type2*)

### 10.3.4 SDS\_GET\_LINK\_TYPE\_PROPERTIES

```
SDS_GET_LINK_TYPE_PROPERTIES (
    sds      : Sds_designator,
    type     : Link_type_nominator_in_sds
)
    category : Category,
    lower_bound : Natural,
    upper_bound : Natural,
    exclusiveness : Exclusiveness,
    stability : Stability,
    duplication : Duplication,
    key_types : Key_types,
    reverse : [ Link_type_nominator_in_sds ]
```

SDS\_GET\_LINK\_TYPE\_PROPERTIES returns the category, lower and upper bounds, exclusiveness, stability, duplication, key attribute types, and reverse link type (if any) of the link type in SDS identified by *type* in the SDS *sds*.

Read locks of the default mode are obtained on *sds* and on the "type" and "type\_in\_sds" objects associated with *type* in *sds*.

**Errors**

ACCESS\_ERRORS (*sds*, ATOMIC, READ, READ\_LINKS)

ACCESS\_ERRORS ("type" and "type\_in\_sds" objects associated with *type*, ATOMIC, READ, READ\_ATTRIBUTES)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *type*)

**10.3.5 SDS\_GET\_OBJECT\_TYPE\_PROPERTIES**

```
SDS_GET_OBJECT_TYPE_PROPERTIES (
    sds           : Sds_designator,
    type          : Object_type_nominator_in_sds
)
    contents_type : [ Contents_type ],
    parents       : Object_type_nominators_in_sds,
    children      : Object_type_nominators_in_sds
```

SDS\_GET\_OBJECT\_TYPE\_PROPERTIES returns the contents type, parents, and children of the object type in SDS identified by *type* in the SDS *sds*.

Read locks of the default mode are obtained on *sds* and on the "type" and "type\_in\_sds" objects associated with *type* in *sds*.

**Errors**

ACCESS\_ERRORS (*sds*, ATOMIC, READ, READ\_LINKS)

ACCESS\_ERRORS ("type" and "type\_in\_sds" objects associated with *type*, ATOMIC, READ, (READ\_ATTRIBUTES, READ\_LINKS))

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *type*)

**10.3.6 SDS\_GET\_TYPE\_KIND**

```
SDS_GET_TYPE_KIND (
    sds           : Sds_designator,
    type          : Type_nominator_in_sds
)
    type_kind    : Type_kind
```

SDS\_GET\_TYPE\_KIND returns the kind of the type in SDS identified by *type* in the SDS *sds*.

Read locks of the default mode are obtained on *sds* and on the "type" and "type\_in\_sds" objects associated with *type* in *sds*.

**Errors**

ACCESS\_ERRORS (*sds*, ATOMIC, READ, READ\_LINKS)

ACCESS\_ERRORS ("type" and "type\_in\_sds" objects associated with *type*, ATOMIC, READ, READ\_ATTRIBUTES)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *type*)

### 10.3.7 SDS\_GET\_TYPE\_MODES

```

SDS_GET_TYPE_MODES (
    sds          : Sds_designator,
    type         : Object_type_nominator_in_sds | Attribute_type_nominator_in_sds |
                  Link_type_nominator_in_sds
)
    usage_mode   : Definition_mode_values,
    export_mode  : Definition_mode_values,
    max_usage_mode : Definition_mode_values

```

SDS\_GET\_TYPE\_MODES returns in *usage\_mode*, *export\_mode*, and *max\_usage\_mode* the usage mode, export mode, and maximum usage mode, respectively, of the type in SDS identified by *type* in the SDS *sds*.

Read locks of the default mode are obtained on *sds* and on the "type" and "type\_in\_sds" objects associated with *type* in *sds*.

#### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, READ, READ\_LINKS)

ACCESS\_ERRORS ("type" and "type\_in\_sds" objects associated with *type*, ATOMIC, READ, READ\_ATTRIBUTES)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *type*)

### 10.3.8 SDS\_GET\_TYPE\_NAME

```

SDS_GET_TYPE_NAME (
    sds      : Sds_designator,
    type     : Type_nominator_in_sds
)
    name : [ Name ]

```

SDS\_GET\_TYPE\_NAME returns the full type name *name* of the type in SDS identified by *type* in the SDS *sds*.

If no name is associated with *type* in *sds* no value is returned.

Read locks of the default mode are obtained on *sds* and on the "type" and "type\_in\_sds" objects associated with *type* in *sds*.

#### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, READ, READ\_LINKS)

ACCESS\_ERRORS ("type" and "type\_in\_sds" objects associated with *type*, ATOMIC, READ, READ\_ATTRIBUTES)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *type*)

### 10.3.9 SDS\_SCAN\_ATTRIBUTE\_TYPE

```

SDS_SCAN_ATTRIBUTE_TYPE (
    sds          : Sds_designator,
    type         : Attribute_type_nominator_in_sds,
    scanning_kind : Attribute_scan_kind
)
    types          : Object_type_nominators_in_sds | Link_type_nominators_in_sds

```

SDS\_SCAN\_ATTRIBUTE\_TYPE returns a set of types *types* determined by *type*, *sds* and *scanning\_kind*.

The returned set of types is determined as follows. It is limited to types associated with types in SDS in the SDS *sds* and by *scanning\_kind* as follows.

- OBJECT: object types to which *type* has been applied by means of SDS\_APPLY\_ATTRIBUTE\_TYPE.
- OBJECT\_ALL: object types of which *type* is an attribute type, i.e. the union of the object types defined by OBJECT and all their descendants.
- LINK\_KEY: link types of which the *type* is a key attribute type.
- LINK\_NON\_KEY: link types of which the *type* is a non-key attribute type.

Read locks of the default mode are obtained on *sds* and on the "type" and "type\_in\_sds" objects associated with the returned types in SDS.

#### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, READ, READ\_LINKS)

ACCESS\_ERRORS ("type" and "type\_in\_sds" objects associated with *type*, ATOMIC, READ, READ\_ATTRIBUTES)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *type*)

### 10.3.10 SDS\_SCAN\_ENUMERAL\_TYPE

```

SDS_SCAN_ENUMERAL_TYPE (
    sds          : Sds_designator,
    type         : Enumeral_type_nominator_in_sds
)
    types          : Attribute_type_nominators_in_sds

```

SDS\_SCAN\_ENUMERAL\_TYPE returns a set of enumeration attribute types, determined by *sds* and the enumerals type *type*.

The returned set of types is limited to types with an associated type in SDS *type\_in\_sds* in the SDS *sds* and to enumeration attribute types with value type containing *type*.

Read locks of the default mode are obtained on *sds* and on the "type" and "type\_in\_sds" objects associated with the returned types in SDS.

#### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, READ, READ\_LINKS)

ACCESS\_ERRORS ("type" and "type\_in\_sds" objects associated with *type*, ATOMIC, READ, READ\_LINKS)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *type*)

The following implementation-defined error may be raised:

ACCESS\_ERRORS ("type" and "type\_in\_sds" objects associated with *types*, ATOMIC, READ, READ\_LINKS)

### 10.3.11 SDS\_SCAN\_LINK\_TYPE

```
SDS_SCAN_LINK_TYPE (
    sds          : Sds_designator,
    type         : Link_type_nominator_in_sds,
    scanning_kind : Link_scan_kind
)
    types       : Object_type_nominators_in_sds | Attribute_type_nominators_in_sds
```

SDS\_SCAN\_LINK\_TYPE returns a set of attribute or object types *types* determined by *sds*, *type*, and *scanning\_kind*.

The returned set of types is determined as follows. It is limited to types with an associated type in SDS *type\_in\_sds* in the SDS *sds* and by *scanning\_kind* as follows.

- ORIGIN: object types which have been defined as origin types of *type* by SDS\_APPLY\_LINK\_TYPE or by SDS\_ADD\_DESTINATION on the reverse link type.
- ORIGIN\_ALL: object types which are valid origins of *type*, i.e. the object types as specified by *scanning\_kind* = ORIGIN, plus all their descendants.
- DESTINATION: object types which have been defined as destination types of *type* by SDS\_ADD\_DESTINATION or by SDS\_APPLY\_LINK\_TYPE on the reverse link type.
- DESTINATION\_ALL: object types which are valid destinations of *type*, i.e. the object types as specified by *scanning\_kind* = DESTINATION plus all their descendants.
- KEY: key attribute types of *type*.
- NON\_KEY: non-key attribute types of *type*, i.e. attribute types which have been applied to *type* by SDS\_APPLY\_ATTRIBUTE\_TYPE.

Read locks of the default mode are obtained on *sds* and on the "type" and "type\_in\_sds" objects associated with the returned types in SDS.

#### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, READ, READ\_LINKS)

ACCESS\_ERRORS ("type" and "type\_in\_sds" objects associated with *type*, ATOMIC, READ, READ\_LINKS)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *type*)

The following implementation-defined error may be raised:

ACCESS\_ERRORS ("type" and "type\_in\_sds" objects associated with *types*, ATOMIC, READ, READ\_LINKS)

### 10.3.12 SDS\_SCAN\_OBJECT\_TYPE

```

SDS_SCAN_OBJECT_TYPE (
    sds          : Sds_designator,
    type         : Object_type_nominator_in_sds,
    scanning_kind : Object_scan_kind
)
    types          : Object_type_nominators_in_sds | Attribute_type_nominators_in_sds |
                    Link_type_nominators_in_sds

```

SDS\_SCAN\_OBJECT\_TYPE returns a set of types *types* determined by *object\_type*, *sds* and *scanning\_kind*.

The returned set of types is determined as follows. It is limited to types with associated types in SDS *type\_in\_sds* in the SDS *sds* and by *scanning\_kind* as follows.

- CHILD: object types which are children of *object\_type*.
- DESCENDANT: object types which are descendants of *object\_type*.
- PARENT: object types which are parents of *object\_type*.
- ANCESTOR: object types which are ancestors of *object\_type*.
- ATTRIBUTE: attribute types which have been applied to *object\_type* by SDS\_APPLY\_ATTRIBUTE\_TYPE.
- ATTRIBUTE\_ALL: attribute types of *object\_type*, i.e. attribute types which have been applied to *object\_type* or to the ancestors of *object\_type*.
- LINK\_ORIGIN: link types which have been applied to *object\_type* (*object\_type* becoming its origin type) by SDS\_APPLY\_LINK\_TYPE or by SDS\_ADD\_DESTINATION on the reverse link type.
- LINK\_ORIGIN\_ALL: link types which have *object\_type* as an origin type, i.e. link types which have been applied to *object\_type* or to its ancestors.
- LINK\_DESTINATION: link types whose destination set has been extended to include *object\_type* by SDS\_ADD\_DESTINATION or by SDS\_APPLY\_LINK\_TYPE on the reverse link type.
- LINK\_DESTINATION\_ALL: link types which have *object\_type* as a destination type, i.e. link types which have been added to the destination set of *object\_type* or to its ancestors.

Read locks of the default mode are obtained on *sds* and on the "type" and "type\_in\_sds" objects associated with the returned types in SDS.

#### Errors

ACCESS\_ERRORS (*sds*, ATOMIC, READ, READ\_LINKS)

ACCESS\_ERRORS ("type" and "type\_in\_sds" objects associated with *type*, ATOMIC, READ, READ\_LINKS)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *object\_type*)

### 10.3.13 SDS\_SCAN\_TYPES

```
SDS_SCAN_TYPES (
    sds      : Sds_designator,
    kind     : [ Type_kind ]
)
    types : Type_nominators_in_sds
```

SDS\_SCAN\_TYPES returns all the types with associated types in SDS in the SDS *sds*, of the kinds given by *kind*.

If *kind* is not supplied, all such types are returned; otherwise all such object types, attribute types, link types, or enumerals are returned according as *kind* is OBJECT\_TYPE, ATTRIBUTE\_TYPE, LINK\_TYPE, or ENUMERAL\_TYPE respectively.

Read locks of the default mode are obtained on *sds* and on the "type" and "type\_in\_sds" objects associated with *type*.

#### Errors

ACCESS\_ERRORS (*sds*, COMPOSITE, READ, READ\_LINKS)

SDS\_IS\_UNKNOWN (*sds*)

## 10.4 Working schema operations

### 10.4.1 WS\_GET\_ATTRIBUTE\_TYPE\_PROPERTIES

```
WS_GET_ATTRIBUTE_TYPE_PROPERTIES (
    type          : Attribute_type_nominator
)
    duplication   : Duplication,
    value_type    : Value_type,
    initial_value : Attribute_value
```

WS\_GET\_ATTRIBUTE\_TYPE\_PROPERTIES returns the duplication, value type identifier, and initial value of the attribute type in working schema associated with the type *type* in the current working schema.

#### Errors

TYPE\_IS\_UNKNOWN\_IN\_WORKING\_SCHEMA (*type*)

### 10.4.2 WS\_GET\_ENUMERAL\_TYPE\_IMAGE

```
WS_GET_ENUMERAL_TYPE_IMAGE (
    type : Enumeral_type_nominator
)
    image : Text
```

WS\_GET\_ENUMERAL\_TYPE\_IMAGE returns the image *image* of the type in working schema associated with the enumerals type *type* in the current working schema.

#### Errors

TYPE\_IS\_UNKNOWN\_IN\_WORKING\_SCHEMA (*type*)



### 10.4.3 WS\_GET\_ENUMERAL\_TYPE\_POSITION

```

WS_GET_ENUMERAL_TYPE_POSITION (
    type1      : Enumeral_type_nominator,
    type2      : Attribute_type_nominator
)
    position   : Natural

```

WS\_GET\_ENUMERAL\_TYPE\_POSITION returns the position *position* of the enumerals type in working schema associated with *type1* in the value type of the attribute type in working schema associated with *type2*, i.e. the key of the "enumerals" link from *type2* to *type1*.

#### Errors

ENUMERAL\_TYPE\_IS\_NOT\_IN\_ATTRIBUTE\_VALUE\_TYPE (*type1*, *type2*)

TYPE\_IS\_UNKNOWN\_IN\_WORKING\_SCHEMA (*type1*)

TYPE\_IS\_UNKNOWN\_IN\_WORKING\_SCHEMA (*type2*)

### 10.4.4 WS\_GET\_LINK\_TYPE\_PROPERTIES

```

WS_GET_LINK_TYPE_PROPERTIES (
    type       : Link_type_nominator
)
    category   : Category,
    lower_bound : Natural,
    upper_bound : Natural,
    exclusiveness : Exclusiveness,
    stability   : Stability,
    duplication : Duplication,
    key_types   : Key_types,
    reverse     : [ Link_type_nominator ]

```

WS\_GET\_LINK\_TYPE\_PROPERTIES returns the category, lower and upper bounds, exclusiveness, stability, duplication, key types, and reverse link type (if any) of the link type in working schema associated with the type *type*.

#### Errors

ACCESS\_ERRORS (*type*, ATOMIC, READ, READ\_ATTRIBUTES)

TYPE\_IS\_UNKNOWN\_IN\_WORKING\_SCHEMA (*type*)

### 10.4.5 WS\_GET\_OBJECT\_TYPE\_PROPERTIES

```

WS_GET_OBJECT_TYPE_PROPERTIES (
    type       : Object_type_nominator
)
    contents_type : [ Contents_type ],
    parents       : Object_type_nominators,
    children      : Object_type_nominators

```

WS\_GET\_OBJECT\_TYPE\_PROPERTIES returns the contents type, parents, and children of the object type in working schema associated with the type *type*.

**Errors**

TYPE\_IS\_UNKNOWN\_IN\_WORKING\_SCHEMA (*type*)

**10.4.6 WS\_GET\_TYPE\_KIND**

```
WS_GET_TYPE_KIND (
    type      : Type_nominator
)
type_kind   : Type_kind
```

WS\_GET\_TYPE\_KIND returns the kind of the type in working schema associated with the type *type* in the current working schema.

**Errors**

TYPE\_IS\_UNKNOWN\_IN\_WORKING\_SCHEMA (*type*)

**10.4.7 WS\_GET\_TYPE\_MODES**

```
WS_GET_TYPE_MODES (
    type      : Type_nominator
)
usage_mode   : Definition_mode_values
```

WS\_GET\_TYPE\_MODES returns the usage mode of the type in working schema associated with the type *type* in the current working schema.

**Errors**

TYPE\_IS\_UNKNOWN\_IN\_WORKING\_SCHEMA (*type*)

**10.4.8 WS\_GET\_TYPE\_NAME**

```
WS_GET_TYPE_NAME (
    type      : Type_nominator
)
name         : [Name]
```

WS\_GET\_TYPE\_NAME returns the first non-null composite name *name* of the type in working schema, considering the sequence of SDSs in the working schema, associated with the type *type* in the current working schema.

If no name is associated with *type* in the current working schema, no value is returned.

**Errors**

TYPE\_IS\_UNKNOWN\_IN\_WORKING\_SCHEMA (*type*)

**10.4.9 WS\_SCAN\_ATTRIBUTE\_TYPE**

```
WS_SCAN_ATTRIBUTE_TYPE (
    type      : Attribute_type_nominator,
    scanning_kind : Attribute_scan_kind
)
types       : Object_type_nominators | Link_type_nominators
```

WS\_SCAN\_ATTRIBUTE\_TYPE returns a set of types *types* determined by *type* and *scanning\_kind*.

The returned set of types is determined as follows. It is limited to types associated with types in working schema in the working schema of the calling process, and by *scanning\_kind* as follows.

- OBJECT: object types to which *type* has been applied by means of SDS\_APPLY\_ATTRIBUTE\_TYPE.
- OBJECT\_ALL: object types of which *type* is an attribute type, i.e. the union of the object types defined by OBJECT and all their descendants.
- LINK\_KEY: link types of which the *type* is a key attribute type.
- LINK\_NON\_KEY: link types of which the *type* is a non-key attribute type.

#### Errors

TYPE\_IS\_UNKNOWN\_IN\_WORKING\_SCHEMA (*type*)

#### 10.4.10 WS\_SCAN\_ENUMERAL\_TYPE

```
WS_SCAN_ENUMERAL_TYPE (
    type : Enumeral_type_nominator
)
    types : Attribute_type_nominators
```

WS\_SCAN\_ENUMERAL\_TYPE returns a set of enumeration attribute types, determined by the enumerals type *type*.

The returned set of types is limited to types with an associated type in working schema in the working schema of the calling process, and to enumeration attribute types with value type containing *type*.

#### Errors

TYPE\_IS\_UNKNOWN\_IN\_WORKING\_SCHEMA (*type*)

#### 10.4.11 WS\_SCAN\_LINK\_TYPE

```
WS_SCAN_LINK_TYPE (
    type : Link_type_nominator,
    scanning_kind : Link_scan_kind
)
    types : Object_type_nominators | Attribute_type_nominators
```

WS\_SCAN\_LINK\_TYPE returns a set of attribute or object types *types* determined by *type* and *scanning\_kind*.

The returned set of types is determined as follows. It is limited to types with an associated type in working schema in the working schema of the calling process, and by *scanning\_kind* as follows.

- ORIGIN: object types which have been defined as origin types of *type* by SDS\_APPLY\_LINK\_TYPE or by SDS\_ADD\_DESTINATION on the reverse link type.

- ORIGIN\_ALL: object types which are valid origins of *type*, i.e. the object types as specified by *scanning\_kind* = ORIGIN plus all their descendants.
- DESTINATION: object types which have been defined as destination types of *type* by SDS\_ADD\_DESTINATION or by SDS\_APPLY\_LINK\_TYPE on the reverse link type.
- DESTINATION\_ALL: object types which are valid destinations of *type*, i.e. the object types as specified by *scanning\_kind* = DESTINATION plus all their descendants.
- KEY: key attribute types of *type*.
- NON\_KEY: non-key attributes of *type*, i.e. attribute types which have been applied to *type* by SDS\_APPLY\_ATTRIBUTE\_TYPE.

### Errors

TYPE\_IS\_UNKNOWN\_IN\_WORKING\_SCHEMA (*type*)

### 10.4.12 WS\_SCAN\_OBJECT\_TYPE

```

WS_SCAN_OBJECT_TYPE (
    object_type      : Object_type_nominator,
    scanning_kind    : Object_scan_kind
)
    types            : Object_type_nominators | Attribute_type_nominators |
                    Link_type_nominators

```

WS\_SCAN\_OBJECT\_TYPE returns a set of types *types* determined by *object\_type* and *scanning\_kind*.

The returned set of types is determined as follows. It is limited to types with associated types in working schema in the working schema of the calling process, and by *scanning\_kind* as follows.

- CHILD: object types which are children of *object\_type*.
- DESCENDANT: object types which are descendants of *object\_type*.
- PARENT: object types which are parents of *object\_type*.
- ANCESTOR: object types which are ancestors of *object\_type*.
- ATTRIBUTE: attribute types which have been applied to *object\_type* by SDS\_APPLY\_ATTRIBUTE\_TYPE.
- ATTRIBUTE\_ALL: attribute types of *object\_type*, i.e. attribute types which have been applied to *object\_type* or to the ancestors of *object\_type*.
- LINK\_ORIGIN: link types which have been applied to *object\_type* (*object\_type* becoming its origin type) by SDS\_APPLY\_LINK\_TYPE or by SDS\_ADD\_DESTINATION on the reverse link type.
- LINK\_ORIGIN\_ALL: link types which have *object\_type* as an origin type, i.e. link types which have been applied to *object\_type* or to its ancestors.
- LINK\_DESTINATION: link types which have had *object\_type* added to their destination object types by SDS\_ADD\_DESTINATION or by SDS\_APPLY\_LINK\_TYPE on the reverse link type.

- LINK\_DESTINATION\_ALL: link types which have had *object\_type* or to its ancestors added to their destination object types.

### Errors

TYPE\_IS\_UNKNOWN\_IN\_WORKING\_SCHEMA (*object\_type*)

### 10.4.13 WS\_SCAN\_TYPES

```
WS_SCAN_TYPES (
    kind    : [ Type_kind ]
)
types     : Type_nominators
```

WS\_SCAN\_TYPES returns all the types of the type kind given by *kind* with associated types in working schema in the current working schema.

If *kind* is not supplied, all such types are returned; otherwise all such object types, attribute types, link types, or enumerals are returned according as *kind* is OBJECT\_TYPE, ATTRIBUTE\_TYPE, LINK\_TYPE, or ENUMERAL\_TYPE respectively.

### Errors

None.

## 11 Volumes, devices, and archives

### 11.1 Volume, device, and archiving concepts

#### 11.1.1 Volumes

Volume\_identifier = Natural

Volume\_accessibility = ACCESSIBLE | INACCESSIBLE | UNKNOWN

Volume\_info = Volume\_identifier \* Volume\_accessibility

Volume\_infos = **set of** Volume\_info

Volume\_status ::

|                   |                     |
|-------------------|---------------------|
| TOTAL_BLOCKS      | : Natural           |
| FREE_BLOCKS       | : Natural           |
| BLOCK_SIZE        | : Natural           |
| NUM_OBJECTS       | : Natural           |
| VOLUME_IDENTIFIER | : Volume_identifier |

**sds** system:

volume\_directory: **child type of** object with **link**

known\_volume: (**navigate**) **non\_duplicated existence link** (volume\_identifier) to volume;

volumes\_of: **implicit link to** common\_root **reverse** volumes;

**end** volume\_directory;

```

volume: child type of object with
attribute
  volume_characteristics: (read) string;
link
  object_on_volume: (navigate) non_duplicated designation link (exact_identifier) to
    object;
  mounted_on: (navigate) non_duplicated designation link to
    device_supporting_volume with
attribute
  read_only: (read) boolean;
end mounted_on;
end volume;
end system;

```

The volume directory is an administrative object (see 9.1.2); it represents the set of known volumes, each with a unique volume identifier which is assigned to the volume on creation and uniquely identifies the volume within the PCTE installation.

The destinations of the "object\_on\_volume" links from a volume are called the objects *residing on* that volume. The value of the "exact\_identifier" attribute is the exact identifier of the object (see 9.1.1). The volume is *mounted* if there is a "mounted\_on" link; the destination of the link is the device that the volume is mounted on (see 11.1.3). The "read\_only" attribute indicates that the volume may not be written to (except for usage designation links, see 8.3.3). A known volume resides on itself; it is the only known volume residing on a volume and it is the first object created on that volume.

The "volume\_characteristics" attribute is an implementation-defined string specifying implementation-dependent characteristics of the volume.

### 11.1.2 Administration volumes

```

sds system:
  administration_volume: (protected) child type of volume with
link
  administration_volume_of: non_duplicated designation link (number) to workstation;
end administration_volume;
end system;

```

Each administration volume is either the master volume or a copy volume of the administration replica set. See 17.1.4.

Each administration volume is associated with one or more workstations (the destinations of the "administration\_volume\_of" links), and is mounted on a device controlled by one of them.

There is exactly one master administration volume in a PCTE installation. It is the master volume of the administration replica set (see 17.1.4), and has volume identifier 0. The master administration volume is part of the initial value of the state (see 8.1).

### 11.1.3 Devices

Device\_identifier = Natural

**sds** system:

device\_supporting\_volume: **child type of device with link**

mounted\_volume: **(navigate) non\_duplicated designation link to** volume;

**end** device\_supporting\_volume;

**end** system;

A device supporting volume is a device (see 12.1) that may have an associated volume, the destination of the "mounted\_volume" link, called the volume *mounted on* the device.

A device supporting volume resides on the administration volume of the workstation controlling it.

#### 11.1.4 Archives

Archive\_selection = Object\_designators | ALL

Archive\_status = PARTIAL | COMPLETE

**sds** system:

archive\_directory: **child type of object with link**

saved\_archive: **(navigate) non\_duplicated existence link** (archive\_identifier: **natural**) **to** archive;

archives\_of: **implicit link to** common root **reverse** archives;

**end** archive\_directory;

archive: **child type of object with attribute**

archiving\_time: **(read) time;**

**link**

archived\_object: **(navigate) non\_duplicated designation link** (exact\_identifier) **to** object;

**end** archive;

**end** system;

The archive directory is an administrative object (see 9.1.2); it represents the set of known archives (the destinations of the "saved\_archive" links), each with a unique archive identifier which is assigned to the archive on creation and uniquely identifies the archive within the PCTE installation.

An archive consists of a set of objects (the destinations of the "archived\_object" links), called the objects *archived on* the archive.

The archiving time of an archive is the system time at which objects are saved in the archive. An archive may only be used once to save objects in it.



## 11.2 Volume, device, and archive operations

### 11.2.1 ARCHIVE\_CREATE

```

ARCHIVE_CREATE (
    archive_identifier      : Natural,
    on_same_volume_as      : Object_designator,
    access_mask            : Atomic_access_rights,
)
    new_archive            : Archive_designator

```

ARCHIVE\_CREATE creates a new archive *new\_archive* residing on the same volume as the object *on\_same\_volume\_as*.

A new "known\_archive" link with key *archive\_identifier* is created from the archive directory to *new\_archive*.

An "object\_on\_volume" link is created from the volume on which *on\_same\_volume\_as* resides to *new\_archive*. The key of the link is the exact identifier of *new\_archive*. *access\_mask* is used in conjunction with the default atomic ACL and default object owner of the calling process to define the atomic ACL and the composite ACL which are to be associated with the created object (see 19.1.4).

The labels of *new\_archive* are set to the mandatory context of the calling process.

Write locks of the default kind are obtained on *new\_archive* and the new "known\_archive" link.

#### Errors

ACCESS\_ERRORS (the archive directory, ATOMIC, MODIFY, APPEND\_LINKS)

ARCHIVE\_EXISTS (*archive\_identifier*)

CONTROL\_WOULD\_NOT\_BE\_GRANTED (*new\_archive*)

LABEL\_IS\_OUTSIDE\_RANGE (*new\_archive*, volume on which *on\_same\_volume\_as* resides)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_CONFIGURATION)

REFERENCE\_CANNOT\_BE\_ALLOCATED

VOLUME\_IS\_FULL (volume on which *on\_same\_volume\_as* resides)

### 11.2.2 ARCHIVE\_REMOVE

```

ARCHIVE_REMOVE (
    archive      : Archive_designator
)

```

ARCHIVE\_REMOVE removes the archive *archive* from the archive directory by deleting the "known\_archive" link to *archive* from the archive directory.

Write locks of the default kind are obtained on *archive* and the deleted "known\_archive" link.

#### Errors

ACCESS\_ERRORS (*archive*, ATOMIC, CHANGE, WRITE\_IMPLICIT)

ACCESS\_ERRORS (*archive*, ATOMIC, MODIFY, DELETE)

ACCESS\_ERRORS (*archive*, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (the archive directory, ATOMIC, MODIFY, WRITE\_LINKS)

ARCHIVE\_HAS\_ARCHIVED\_OBJECTS (*archive*)

ARCHIVE\_IS\_UNKNOWN (*archive*)

OBJECT\_IS\_IN\_USE\_FOR\_DELETE (*archive*)

OBJECT\_IS\_INACCESSIBLE (*archive*, ATOMIC)

If the conditions hold for deletion of the "archive" object *archive*:

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_CONFIGURATION)

### 11.2.3 ARCHIVE\_RESTORE

```

ARCHIVE_RESTORE (
    device           : Device_designator,
    archive          : Archive_designator,
    scope            : Archive_selection,
    on_same_volume_as : Object_designator
)
    restoring_status : Archive_status

```

ARCHIVE\_RESTORE restores a set of objects *objects* specified by *scope* to the volume *volume* on which *on\_same\_volume\_as* resides from the archive *archive*.

If *scope* is a set of object designators, the specified set of objects to be restored (called the 'specified set' in this clause) is the intersection of the set of objects archived on *archive* and the set of objects in *scope*.

If *scope* is ALL, the specified set is the set of all the objects archived on *archive*.

The objects to be restored are taken from the contents of *device*.

The objects and their components are moved to *volume*, in an undefined order; as many objects and their components as possible are restored from *device*, and reside on *volume*.

If not all the objects of the specified set can be restored by this operation, as many as possible are restored and *restoring\_status* is set to PARTIAL. If all the objects of the specified set are restored, *restoring\_status* is set to COMPLETE. If no objects of the specified set can be restored, the error condition VOLUME\_IS\_FULL is raised.

The "archived\_object" links from *archive* to the restored objects are deleted.

For each of the objects which are restored to *volume*, an "object\_on\_volume" link with key the exact identifier of the object is created.

If any of the objects specified to be restored has not been archived or is already restored on *volume*, then it is not affected.

Write locks of the default mode are obtained on *archive* and on the moved objects and links. A read lock of the default mode is obtained on *device*.

#### Errors

ACCESS\_ERRORS (*device*, ATOMIC, READ, READ\_CONTENTS)

ACCESS\_ERRORS (an element of *scope*, COMPOSITE, CHANGE, CONTROL\_OBJECT)

ARCHIVE\_IS\_INVALID\_ON\_DEVICE (*device*, *archive*)

LABEL\_IS\_OUTSIDE\_RANGE (an element of the specified set or a component of such an element, *volume*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_CONFIGURATION)

PROCESS\_IS\_IN\_TRANSACTION

VOLUME\_IS\_FULL (*volume*)

VOLUME\_IS\_INACCESSIBLE (*volume*)

VOLUME\_IS\_READ\_ONLY (*scope*, COMPOSITE)

The following implementation-dependent errors may be raised for any object X with a link to an object of *objects*:

OBJECT\_IS\_INACCESSIBLY\_ARCHIVED (X)

VOLUME\_IS\_INACCESSIBLE (volume on which X resides)

VOLUME\_IS\_READ\_ONLY (volume on which X resides)

#### 11.2.4 ARCHIVE\_SAVE

```

ARCHIVE_SAVE (
    device          : Device_designator,
    archive         : Archive_designator,
    objects         : Object_designators
)
    archiving_status : Archive_status
  
```

ARCHIVE\_SAVE moves a set of objects to the contents of *device*.

The archive *archive* is updated as follows:

- the archiving time is set to the current system time;
- "archived\_object" links are created from *archive* to the archived objects and to each of their components. The keys of the created links are the suffixes of the exact identifier of the destination objects.

For each archived object, the "object\_on\_volume" link from the volume on which the object resides to the object is deleted.

If *device* has insufficient space to hold all the objects of *objects* and their components, the operation archives as many objects as possible and *archiving\_status* is set to PARTIAL. If *device* has insufficient space to hold any objects of *objects* with their components, the error condition DEVICE\_SPACE\_IS\_FULL occurs. Otherwise *archiving\_status* is set to COMPLETE.

The operation has no effect on objects or components which are already archived, either on the same archive or on another one.

Read locks of the default mode are obtained on the objects to be archived. Write locks of the default mode are obtained on *archive* and on *device*.

#### Errors

ACCESS\_ERRORS (*device*, ATOMIC, MODIFY, WRITE\_CONTENTS)

ACCESS\_ERRORS (*archive*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (elements of *objects* and their components that are to be archived, ATOMIC, CHANGE, CONTROL\_OBJECT)

ARCHIVE\_HAS\_ARCHIVED\_OBJECTS (*archive*)

DEVICE\_SPACE\_IS\_FULL (*device*)

LABEL\_IS\_OUTSIDE\_RANGE (an element or a component of an element of *objects*, *device*)

OBJECT\_ARCHIVING\_IS\_INVALID (*objects*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_CONFIGURATION)

PROCESS\_IS\_IN\_TRANSACTION

The following implementation-dependent errors may be raised for any object X with a link to an object of *objects*:

OBJECT\_IS\_INACCESSIBLY\_ARCHIVED (X)

VOLUME\_IS\_INACCESSIBLE (volume on which X resides)

VOLUME\_IS\_READ\_ONLY (volume on which X resides)

NOTE - It is intended that the space previously occupied by the archived objects be freed.

### 11.2.5 DEVICE\_CREATE

```

DEVICE_CREATE (
    station           : Workstation_designator,
    device_type       : Device_type_nominator,
    access_mask       : Atomic_access_rights,
    device_identifier  : Natural,
    device_characteristics : String
)
    new_device       : Device_designator
  
```

DEVICE\_CREATE creates a device *new\_device* of type *device\_type* with a "controlled\_device" link to it from *station*. The value of *device\_identifier* is the key of the created link. The "device\_of" reverse link created from the new object to *station* designates the workstation which controls the device. The "device\_characteristics" attribute of *new\_device* is set to *device\_characteristics*.

*device\_identifier* is a value which uniquely identifies the new device within the devices controlled by *station*. Its value is assigned to the "device\_identifier" attribute of *new\_device*.

*new\_device* resides on the same volume as *station* (i.e. the local administration volume of *station*) and cannot be moved to another volume.

*access\_mask* is used in conjunction with the default atomic ACL and default object owner of the calling process to define both the atomic ACL and the composite ACL which are to be associated with the created object (see 19.1.4).

An "object\_on\_volume" link is created from the administration volume of *station* to *new\_device*. The created link is keyed by the exact identifier of *new\_device*.

The security labels of *new\_device* and the labels defining its security ranges are set to the mandatory context of the calling process.

Write locks (of the default kind) are obtained on *new\_device* and on the new links (except the new "object\_on\_volume" link).

#### Errors

ACCESS\_ERRORS (*station*, ATOMIC, MODIFY, APPEND\_LINKS)

CONTROL\_WOULD\_NOT\_BE\_GRANTED (*new\_device*)

DEVICE\_CHARACTERISTICS\_ARE\_INVALID (*device-characteristics*)  
 DEVICE\_EXISTS (*device\_identifier*)  
 LABEL\_IS\_OUTSIDE\_RANGE (*new\_device*, *station*)  
 LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_KEY\_VALUE)  
 OBJECT\_OWNER\_VALUE\_WOULD\_BE\_INCONSISTENT\_WITH\_ATOMIC\_ACL  
 PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_CONFIGURATION)  
 REFERENCE\_CANNOT\_BE\_ALLOCATED  
 USAGE\_MODE\_ON\_OBJECT\_TYPE\_WOULD\_BE\_VIOLATED ("object", *device\_type*)  
 WORKSTATION\_IS\_UNKNOWN (*station*)

### 11.2.6 DEVICE\_REMOVE

```

    DEVICE_REMOVE (
        device : Device_designator
    )
  
```

DEVICE\_REMOVE removes the device object *device* from the set of devices of a workstation. As a result, the device object *device* does not represent a physical device and its associated device identifier can be reused.

The "controlled\_device" link from the workstation to *device* is deleted. If it is the only existence link to *device* and there are no composition links to *device*, *device* is also deleted. In that case, the "object\_on\_volume" link from the volume on which *device* was residing to *device* is also deleted.

A write lock (of the default kind) is obtained on *device* if it is deleted and on the deleted links (except the "object\_on\_volume" link).

#### Errors

ACCESS\_ERRORS (*device*, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (*station*, ATOMIC, MODIFY, WRITE\_LINKS)

If conditions hold for the deletion of *device*:

ACCESS\_ERRORS (*device*, COMPOSITE, MODIFY, DELETE)

For each origin X of an implicit link to *device*:

ACCESS\_ERRORS (X, ATOMIC, CHANGE, WRITE\_IMPLICIT)

For each atomically stabilizing link L of *device*:

ACCESS\_ERRORS (destination of L, ATOMIC, CHANGE, STABILIZE)

For each compositely stabilizing link L of *device*:

ACCESS\_ERRORS (destination of L, COMPOSITE, CHANGE, STABILIZE)

DEVICE\_IS\_IN\_USE (*device*)

DEVICE\_IS\_UNKNOWN (*device*)

If conditions hold for the deletion of *device*:

OBJECT\_HAS\_LINKS\_PREVENTING\_DELETION (*device*)

OBJECT\_IS\_IN\_USE\_FOR\_DELETE (*device*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_CONFIGURATION)

NOTE - This operation prevents further use of the device.

### 11.2.7 LINK\_GET\_DESTINATION\_ARCHIVE

```

LINK_GET_DESTINATION_ARCHIVE (
    origin           : Object_designator,
    link            : Link_designator
)
    archive_identifier : Archive_identifier

```

LINK\_GET\_DESTINATION\_ARCHIVE returns the archive identifier of the destination object of the direct outgoing link *link* of the object *origin*.

A read lock of default mode is obtained on *link*.

#### Errors

ACCESS\_ERRORS (*origin*, ATOMIC, READ, READ\_LINKS)

OBJECT\_IS\_NOT\_ARCHIVED (destination object of *link*)

### 11.2.8 VOLUME\_CREATE

```

VOLUME_CREATE (
    device           : Device_supporting_volume_designator,
    volume_identifier : Natural,
    access_mask      : Atomic_access_rights,
    volume_characteristics : String
)
    new_volume       : Volume_designator

```

VOLUME\_CREATE creates a new volume *new\_volume* and mounts it on the device *device*. *new\_volume* resides on itself.

A new "known\_volume" link with key *volume\_identifier* is created from the master of the volume directory to *new\_volume*.

A "mounted\_on" link with "read\_only" attribute set to **false** and its reverse are created between *new\_volume* and *device*.

An "object\_on\_volume" link is created from *new\_volume* to itself. The key of the link is the exact identifier of *new\_volume*. *access\_mask* is used in conjunction with the default atomic ACL and default object owner of the calling process to define the atomic ACL and the composite ACL which are to be associated with the created object (see 19.1.4).

The labels of the volume and the labels defining its security ranges are set to the mandatory context of the calling process. Each security range of the created volume must lie within the corresponding security range of *device* (see 20.1.5).

The "volume\_characteristics" attribute of *new\_volume* is set to *volume\_characteristics*.

Write locks of the default mode are obtained on *new\_volume* and the new links (except the new "object\_on\_volume" link).

#### Errors

ACCESS\_ERRORS (*device*, ATOMIC, MODIFY, EXPLOIT\_DEVICE)

ACCESS\_ERRORS (the directory of volumes, ATOMIC, MODIFY, APPEND\_LINKS)

CONTROL\_WOULD\_NOT\_BE\_GRANTED (*new\_version*)

DEVICE\_IS\_BUSY (*device*, *volume\_identifier*)  
 DEVICE\_IS\_UNKNOWN (*device*)  
 LABEL\_IS\_OUTSIDE\_RANGE (*new\_volume*, *device*)  
 LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_KEY\_VALUE)  
 OBJECT\_OWNER\_VALUE\_WOULD\_BE\_INCONSISTENT\_WITH\_ATOMIC\_ACL  
 PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_CONFIGURATION)  
 PROCESS\_IS\_IN\_TRANSACTION  
 REFERENCE\_CANNOT\_BE\_ALLOCATED  
 VOLUME\_EXISTS (*volume\_identifier*)

NOTE - The new volume may need to be initialized by a system tool before this operation is called.

### 11.2.9 VOLUME\_DELETE

```

VOLUME_DELETE (
    volume : Volume_designator
)
  
```

VOLUME\_DELETE unmounts the volume *volume*, and deletes the "known\_volume" link to *volume* from the master of the volume directory and the "mounted\_volume" link from the device on which *volume* is mounted.

*volume* must be the only object residing on *volume*, and there must be only the following three links from *volume*:

- the reverse link of the "known\_volume" link to *volume* from the volume directory;
- the "object\_on\_volume" link from *volume* to itself;
- the "mounted\_on" link from *volume*.

Write locks (of the default kind) are obtained on *volume* and the deleted links (except the "object\_on\_volume" link); however the locks on *volume* and on the links from *volume* do not prevent the unmounting of the volume.

#### Errors

ACCESS\_ERRORS (*device*, ATOMIC, MODIFY, EXPLOIT\_DEVICE)  
 ACCESS\_ERRORS (the volume directory, ATOMIC, MODIFY, WRITE\_LINKS)  
 ACCESS\_ERRORS (*volume*, ATOMIC, MODIFY, WRITE\_LINKS)  
 ACCESS\_ERRORS (*volume*, ATOMIC, CHANGE, WRITE\_IMPLICIT)  
 If the conditions hold for deletion of the "volume" object *volume*:  
     ACCESS\_ERRORS (*volume*, ATOMIC, MODIFY, DELETE)  
 PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_CONFIGURATION)  
 PROCESS\_IS\_IN\_TRANSACTION  
 VOLUME\_HAS\_OTHER\_LINKS (*volume*)  
 VOLUME\_HAS\_OTHER\_OBJECTS (*volume*)  
 VOLUME\_IS\_INACCESSIBLE (*volume*)  
 VOLUME\_IS\_UNKNOWN (*volume*)



### 11.2.10 VOLUME\_GET\_STATUS

```
VOLUME_GET_STATUS (
    volume      : Volume_designator
)
    status      : Volume_status
```

VOLUME\_GET\_STATUS returns information about the mounted volume *volume*, as follows.

- TOTAL\_BLOCKS is the total number of blocks of data available on *volume*.
- FREE\_BLOCKS is the number of blocks of data which are free on *volume*.
- BLOCK\_SIZE is the size of a block of data on *volume*, in octets.
- NUM\_OBJECTS is the number of objects currently residing on *volume*.
- VOLUME\_IDENTIFIER is the volume identifier of *volume*.

A read lock of the default mode is obtained on *volume*.

#### Errors

ACCESS\_ERRORS (*volume*, ATOMIC, READ, READ\_ATTRIBUTES)

VOLUME\_IS\_INACCESSIBLE (*volume*)

VOLUME\_IS\_UNKNOWN (*volume*)

### 11.2.11 VOLUME\_MOUNT

```
VOLUME_MOUNT (
    device      : Device_supporting_volume_designator,
    volume_identifier : Volume_identifier,
    read_only    : Boolean
)
```

VOLUME\_MOUNT causes the volume *volume* identified by *volume\_identifier* to be mounted on the device *device*.

The operation creates a "mounted\_on" link from *volume* to *device*, with "read\_only" attribute set to *read\_only*, and its reverse "mounted\_volume" link.

A lock of external and internal mode READ\_SEMIPROTECTED is established on *device*.

Write locks (of the default kind) are obtained on the created links.

#### Errors

ACCESS\_ERRORS (*device*, ATOMIC, MODIFY, EXPLOIT\_DEVICE)

ACCESS\_ERRORS (*volume*, ATOMIC, READ, NAVIGATE)

DEVICE\_IS\_BUSY (*device*)

DEVICE\_IS\_UNKNOWN (*device*)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_MOUNTED\_VOLUMES)

PROCESS\_IS\_IN\_TRANSACTION

RANGE\_IS\_OUTSIDE\_RANGE (*volume* associated with *volume*, *device*)

VOLUME\_CANNOT\_BE\_MOUNTED\_ON\_DEVICE (*volume*, *device*)

VOLUME\_IS\_ALREADY\_MOUNTED (*volume*)

VOLUME\_IS\_UNKNOWN (*volume*)

NOTE - When appropriate, the operation causes the physical mounting of the corresponding physical volume on the corresponding physical device. If *read\_only* is true then the physical volume is mounted for reading only.

### 11.2.12 VOLUME\_UNMOUNT

```
VOLUME_UNMOUNT (
    volume : Volume_designator
)
```

VOLUME\_UNMOUNT causes the volume *volume* to be unmounted.

The "mounted\_on" link from *volume* to the device *device* on which *volume* is mounted is deleted.

Write locks (of the default kind) are obtained on the volume and on the deleted link; however the locks on the volume and on the link from that object do not prevent the unmounting of the volume.

#### Errors

ACCESS\_ERRORS (*device*, ATOMIC, MODIFY, EXPLOIT\_DEVICE)

ACCESS\_ERRORS (*volume*, ATOMIC, READ, NAVIGATE)

PROCESS\_IS\_IN\_TRANSACTION

VOLUME\_HAS\_OBJECTS\_IN\_USE (*volume*)

VOLUME\_IS\_ADMINISTRATION\_VOLUME (*volume*)

VOLUME\_IS\_INACCESSIBLE (*volume*)

VOLUME\_IS\_UNKNOWN (*volume*)

## 12 Files, pipes, and devices

### 12.1 File, pipe, and device concepts

Open\_contents ::

OPEN\_OBJECT\_KEY : Natural

CURRENT\_POSITION: Current\_position

Current\_position :: Token

Contents\_handle :: Token

Position\_handle :: Token

Contents\_access\_mode = READ\_WRITE | READ\_ONLY | WRITE\_ONLY | APPEND\_ONLY

Seek\_position = FROM\_BEGINNING | FROM\_CURRENT | FROM\_END

Set\_position = AT\_BEGINNING | AT\_POSITION | AT\_END

File = seq of Octet

represented by file

Pipe = seq of Octet

represented by pipe

Device = **seq** of Octet  
**represented by** device

Control\_data = **seq** of Octet

Positioning\_style = SEQUENTIAL | DIRECT | SEEK

**sds** system:

positioning: (**read**) **enumeration** (SEQUENTIAL, DIRECT, SEEK) := SEQUENTIAL;

file: **child type of** object with  
**contents file**;  
**attribute**  
    contents\_size: (**read**) **natural**;  
    positioning;  
**end file**;

pipe: **child type of** object with  
**contents pipe**;  
**end pipe**;

device: **child type of** object with  
**contents device**;  
**attribute**  
    device\_characteristics: (**read**) **string**;  
    positioning;  
**link**  
    device\_of: (**navigate**) **reference link to** workstation **reverse** controlled\_device;  
**end device**;

**end system**;

The contents of a file, pipe, or device may be accessed by a process as a sequence of octets if it is *opened* by the process. The file, pipe, or device is then called an *open object*.

The opening of a file, pipe, or device by a process is represented by an "open\_object" link from the process to the file, pipe, or device, and an "opened\_by" link from the file, pipe or device to the process.

The "open\_object\_key" key attribute of the "open\_object" link uniquely identifies the open object among the other objects opened by the process.

The opening of an object's contents and the operation CONTENTS\_GET\_HANDLE\_FROM\_KEY result in the creation of a *contents handle*, which is an implementation-dependent reference to the open contents of an object. A separate open contents value exists for each process that opens a particular contents. The open contents consists of the following:

- An *open object key* which is the key of an "open\_object" link to the open object.
- A *current position* which specifies one octet of the sequence within the logical sequence of octets. The position of the first octet is called FIRST, and of the last octet is called LAST. A current position can be shared by several open contents.

The "device\_characteristics" attribute of a device is a string with an implementation-defined syntax specifying implementation-dependent characteristics of the device.

The "positioning" attribute of files and devices can be set only by CONTENTS\_SET\_PROPERTIES. It is defined as follows:

- SEQUENTIAL indicates that the current position can be changed only by writing or reading octets in a sequential way.
- DIRECT indicates that the current position can be changed either as by SEQUENTIAL or by means of a previously saved position, represented by an implementation-dependent *position handle*.
- SEEK indicates that the current position can be changed either as by DIRECT or by an offset from another position.

The contents of a pipe is always accessed sequentially.

The contents size of a file is 0 if the file is empty and otherwise  $LAST - FIRST + 1$ .

The "open\_object" link has an "opening\_mode" attribute which defines how the current position is updated by the contents operations. Let CP be the current position, and CP+n the position of the *n*th octet after the current position. An operation is allowed for any opening mode unless otherwise stated below.

In READ\_WRITE opening mode:

- opening the contents by CONTENTS\_OPEN sets CP to FIRST;
- successfully reading N octets by CONTENTS\_READ returns the octets at positions CP, CP+1, ... CP+N-1, and changes CP to CP+N;
- successfully writing N octets by CONTENTS\_WRITE replaces or adds octets at positions CP, CP+1, ... CP+N-1, and changes CP to CP+N; and if  $LAST < CP+N$ , changes LAST to CP+N.

In READ\_ONLY opening mode:

- opening the contents with CONTENTS\_OPEN sets CP to FIRST;
- successfully reading N octets by CONTENTS\_READ returns the octets at positions CP, CP+1, ... CP+N-1, and changes CP to CP+N. For pipes and *read-once devices*, e.g. keyboards, the sequence of octets is changed to identify as FIRST the new current position. Which devices are read-once is implementation-defined.
- CONTENTS\_WRITE and CONTENTS\_TRUNCATE are not allowed.

In WRITE\_ONLY opening mode:

- opening the contents by CONTENTS\_OPEN sets CP to FIRST;
- CONTENTS\_READ is not allowed;
- successfully writing N octets by CONTENTS\_WRITE replaces or adds octets at positions CP, CP+1, ... CP+N-1, and changes CP to CP+N; and if  $LAST < CP+N$ , changes LAST to CP+N.

In APPEND\_ONLY opening mode:

- opening the contents by CONTENTS\_OPEN sets CP to LAST+1;
- CONTENTS\_READ is not allowed;
- successfully writing N octets by CONTENTS\_WRITE sets CP to LAST+1, adds octets at positions LAST+1, LAST+2, ..., LAST+N, and changes LAST to LAST+N;
- CONTENTS\_SET\_POSITION, CONTENTS\_SEEK, and CONTENTS\_TRUNCATE are not allowed.

In READ\_ONLY, WRITE\_ONLY and READ\_WRITE opening modes:

- if positioning is DIRECT or SEEK, positioning with CONTENTS\_SET\_POSITION sets CP to a position identified by a position handle;
- if positioning is SEEK, positioning with CONTENTS\_SEEK sets CP to a position identified by an offset from another position.

Pipes can be opened only in READ\_ONLY or APPEND\_ONLY mode.

The "open\_object" link has an "inheritable" attribute; if it is **true**, then a link of the same type, key, non-key system attributes, and destination object is created from any process created by the process origin of the link.

The "open\_object" link has a "non\_blocking\_io" attribute which defines the behaviour of the operations CONTENTS\_READ and CONTENTS\_WRITE. This property is always **true** for a file, but is **true** for a pipe or a device only if it supports non-blocking input-output, i.e. a read or write operation does not wait until all data can be read or written, but reads or writes as much as it can. If it is **true**, then the destination pipe or device is said to be *non-blocking* (for the opening process).

If an "open\_object" link is inheritable, the associated contents handle is duplicated in the child process. As a consequence, the parent and child process share the current position.

The effect of changing the current position by the operations CONTENTS\_READ, CONTENTS\_WRITE, CONTENTS\_TRUNCATE, CONTENTS\_SEEK, and CONTENTS\_SET\_POSITION is visible to all processes sharing that current position.

#### NOTES

1 The "open\_object" links keyed by 0, 1, and 2 are called *standard input*, *standard output*, and *standard error* respectively. Conventionally, standard input is used by the process for the reading of commands or input data, standard output is used for the output of data, and standard error is used for the output of error diagnostics.

2 After creating a process, an "opened\_object" link is created for each "open\_object" link which has the inheritable property set to true. The designated object is however not open until the process is started: the starting of the process creates a link of type "open\_by" from the designated object to the process.

3 In APPEND\_ONLY opening mode, the current position is always LAST + 1. Several processes can append to the same file, pipe, or device and therefore concurrently modify the LAST position subject to locking rules.

4 On pipes and some kinds of devices (e.g. keyboards), the reading of a sequence of octets deletes the octets: the sequence of octets read is no longer readable and the new current position identifies as FIRST the next unread octet in the sequence. Several processes can concurrently read the same pipe or device in this way.

5 There are various situations allowing one or more contents handles to be associated with the same object in such a way that the CONTENTS\_READ, CONTENTS\_WRITE and CONTENTS\_TRUNCATE operations performed on the two contents handles may interfere:

- two contents handles opened within the context of the same activity, either within the same process or within different processes;
- contents handles obtained from objects locked within concurrent activities but with compatible locks.

6 An application needing to manage such interferences without using separate activities and appropriate locks must use its own synchronization mechanisms.

7 The contents of pipes and devices are not affected by transaction rollback.

## 12.2 File, pipe, and device operations

### 12.2.1 CONTENTS\_CLOSE

```

CONTENTS_CLOSE (
    contents : Contents_handle
)

```

CONTENTS\_CLOSE deletes the contents handle contents, releasing any associated resources.

The "open\_object" link keyed by the open object key of *contents* and its complementary "opened\_by" link (see 12.2.6) are deleted.

#### Errors

CONTENTS\_IS\_NOT\_OPEN (*contents*)

### 12.2.2 CONTENTS\_GET\_HANDLE\_FROM\_KEY

```

CONTENTS_GET_HANDLE_FROM_KEY (
    open_object_key : Natural
)
    contents : Contents_handle

```

CONTENTS\_GET\_HANDLE\_FROM\_KEY returns in *contents* the contents handle of the calling process represented by the "open\_object" link with key *open\_object\_key*.

#### Errors

LINK\_DOES\_NOT\_EXIST (calling process, *link*)

### 12.2.3 CONTENTS\_GET\_KEY\_FROM\_HANDLE

```

CONTENTS_GET_KEY_FROM_HANDLE (
    contents : Contents_handle
)
    open_object_key : Natural

```

CONTENTS\_GET\_KEY\_FROM\_HANDLE returns in *open\_object\_key* the "open\_object\_key" key attribute of the "open\_object" link associated with the contents handle *contents*.

#### Errors

CONTENTS\_IS\_NOT\_OPEN (*contents*)

### 12.2.4 CONTENTS\_GET\_POSITION

```

CONTENTS_GET_POSITION (
    contents : Contents_handle
)
    position : Position_handle

```

CONTENTS\_GET\_POSITION returns in *position* a position handle representing the current position of *contents*.

**Errors**

CONTENTS\_IS\_NOT\_OPEN (*contents*)

If *contents* is a pipe, or if *contents* is a file or a device with positioning SEQUENTIAL:

CONTENTS\_OPERATION\_IS\_INVALID (*contents*)

**12.2.5 CONTENTS\_HANDLE\_DUPLICATE**

```

CONTENTS_HANDLE_DUPLICATE (
    contents      : Contents_handle,
    new_key       : [ Natural ],
    inheritable   : Boolean
)
    new_contents : Contents_handle

```

CONTENTS\_HANDLE\_DUPLICATE creates a new open contents for the calling process and the object associated with *contents*, and returns a contents handle identifying it in *new\_contents*. A new "open\_object" link from the calling process to the object identified by *contents*, and a complementary "opened\_by" link, are created.

The "inheritable" attribute of the new "open\_contents" link is set to *inheritable*. The key of the new "open\_object" link is set to *new\_key*, if provided, and otherwise to an unused implementation-defined value. The "opening\_mode" and "non\_blocking\_io" attributes of the new "open\_object" link are set to the same values as for the "open\_object" link from the calling process associated with *contents*.

The new open contents shares the current position of *contents*.

**Errors**

CONTENTS\_IS\_NOT\_OPEN (*contents*)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_OPEN\_OBJECTS)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_OPEN\_OBJECTS\_PER\_PROCESS)

OPEN\_KEY\_IS\_INVALID (*new\_key*)

**12.2.6 CONTENTS\_OPEN**

```

CONTENTS_OPEN (
    object          : File_designator | Pipe_designator | Device_designator,
    opening_mode    : Contents_access_mode,
    non_blocking_io : Boolean,
    inheritable     : Boolean
)
    contents        : Contents_handle

```

CONTENTS\_OPEN opens the contents of *object* in the opening mode *opening\_mode* and returns a contents handle for it in *contents*.

An "open\_object" link is created from the calling process to *object* with opening mode set to *opening\_mode*, non-blocking io set to *non\_blocking\_io*, and inheritable set to *inheritable*.

An "opened\_by" link is created with an implementation-dependent key from *object* to the calling process, complementary to the created "open\_object" link.



If *opening\_mode* is READ\_ONLY, a read lock of the default mode is obtained on *object*.

If *opening\_mode* is WRITE\_ONLY, READ\_WRITE or APPEND\_ONLY, a write lock of the default mode is obtained on *object*.

After this operation, the object is operated on by the current activity and the lock established is not released before the contents is closed (see 16.1.8).

### Errors

If *opening\_mode* is READ\_ONLY or READ\_WRITE:

ACCESS\_ERRORS (*object*, ATOMIC, READ, READ\_CONTENTS)

If *opening\_mode* is WRITE\_ONLY or READ\_WRITE:

ACCESS\_ERRORS (*object*, ATOMIC, MODIFY, WRITE\_CONTENTS)

If *opening\_mode* is APPEND\_ONLY

ACCESS\_ERRORS (*object*, ATOMIC, MODIFY, APPEND\_CONTENTS)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_OPEN\_OBJECTS)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_OPEN\_OBJECTS\_PER\_PROCESS)

NON\_BLOCKING\_IO\_IS\_INVALID (*object*, *non\_blocking\_io*)

OPENING\_MODE\_IS\_INVALID (*object*, *opening\_mode*)

STATIC\_CONTEXT\_IS\_IN\_USE (*object*)

### 12.2.7 CONTENTS\_READ

```

CONTENTS_READ (
    contents : Contents_handle,
    size      : Natural
)
    data      : Unstructured_contents
  
```

CONTENTS\_READ reads a sequence of *size* octets from *contents* at the current position and returns it in *data*, if available. If there are less than *size* octets but at least one octet from the current position to LAST inclusive, the operation returns in *data* that sequence of octets.

The current position is set to the position after the last read octet.

If *contents* is a pipe or a read-once device, the position after the last read octet is identified as FIRST after the operation.

If there are no octets available for reading:

- if *contents* is a non-blocking pipe or device, the operation fails.
- if *contents* is a blocking device or a blocking pipe for which a contents handle is open in APPEND\_ONLY mode, then the operation waits until some octets are available for reading. If, in the case of a pipe, the last contents handle open in APPEND\_ONLY mode is closed while the operation is waiting, an empty sequence of octets is returned.
- if *contents* is a file, *data* is set to the empty sequence.

### Errors

CONFIDENTIALITY\_WOULD\_BE\_VIOLATED (*object* determined by *contents*, ATOMIC)

CONTENTS\_IS\_NOT\_OPEN (*contents*)

CONTENTS\_OPERATION\_IS\_INVALID (*contents*)  
 DATA\_ARE\_NOT\_AVAILABLE (*contents*)  
 INTEGRITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED (object determined by *contents*, ATOMIC)  
 OBJECT\_IS\_INACCESSIBLE (object determined by *contents*, ATOMIC)  
 PIPE\_HAS\_NO\_WRITERS (*contents*)  
 VOLUME\_IS\_INACCESSIBLE (volume on which object determined by *contents* resides, ATOMIC)

### 12.2.8 CONTENTS\_SEEK

```

CONTENTS_SEEK (
    contents      : Contents_handle,
    offset        : Integer,
    whence        : Seek_position
)
    new_position : Natural
  
```

CONTENTS\_SEEK sets the current position of *contents* to a position determined by *offset* and *whence*, and returns the new current position as an offset from FIRST.

If *contents* is a file or device which has the SEEK "positioning" property, the current position is set to a position defined as follows:

- if *whence* is FROM\_BEGINNING, the new position is FIRST + *offset*;
- if *whence* is FROM\_CURRENT, the new position is the current position + *offset*;
- if *whence* is FROM\_END, the new position is LAST + *offset* + 1.

The resulting position cannot be smaller than FIRST.

The resulting position RP can be greater than LAST. In this case, if subsequent writing occurs, LAST and the "contents\_size" attribute of the file are set to RP + number of written octets. The octets which are between the previous LAST position and RP are returned as octets with the value 0 by subsequent calls of CONTENTS\_READ. However, the file remains unchanged if no CONTENTS\_WRITE occurs at the new RP position.

If *contents* is a file or a device, the new value of the current position, offset from the beginning of the file, is returned in *new\_position*.

#### Errors

CONTENTS\_IS\_NOT\_OPEN (*contents*)  
 CONTENTS\_OPERATION\_IS\_INVALID (*contents*)  
 OBJECT\_IS\_INACCESSIBLE (object determined by *contents*, ATOMIC)  
 POSITION\_IS\_INVALID (resulting position)

### 12.2.9 CONTENTS\_SET\_POSITION

```

CONTENTS_SET_POSITION (
    contents      : Contents_handle,
    position_handle : Position_handle,
    set_mode      : Set_position
)

```

CONTENTS\_SET\_POSITION sets the current position of *contents* to a position determined by *set\_mode* and *position\_handle*.

If *set\_mode* is AT\_BEGINNING or AT\_END, the current position of *contents* is set to FIRST or LAST + 1 respectively.

If *set\_mode* is AT\_POSITION, the current position of *contents* is set to the position represented by *position\_handle* which must have been previously obtained by a call of CONTENTS\_GET\_POSITION on *contents*.

#### Errors

CONTENTS\_IS\_NOT\_OPEN (*contents*)

CONTENTS\_OPERATION\_IS\_INVALID (*contents*)

OBJECT\_IS\_INACCESSIBLE (object determined by *contents*, ATOMIC)

POSITION\_HANDLE\_IS\_INVALID (*position\_handle*, *contents*)

### 12.2.10 CONTENTS\_SET\_PROPERTIES

```

CONTENTS_SET_PROPERTIES (
    contents      : Contents_handle,
    positioning    : Positioning_style
)

```

CONTENTS\_SET\_PROPERTIES sets the positioning of the open file or device determined by *contents* to *positioning*.

If *contents* determines a file, its positioning can be changed only if the file is empty.

#### Errors

If *contents* determines a file:

CONTENTS\_IS\_NOT\_EMPTY (*contents*)

CONTENTS\_IS\_NOT\_OPEN (*contents*)

If *contents* determines a pipe, or a file or device open in mode READ\_ONLY:

CONTENTS\_OPERATION\_IS\_INVALID (*contents*)

OBJECT\_IS\_INACCESSIBLE (object determined by *contents*, ATOMIC)

POSITIONING\_IS\_INVALID (*contents*, *positioning*)

NOTE – The change of properties is made on behalf of the activity in which the contents was opened.

### 12.2.11 CONTENTS\_TRUNCATE

```

CONTENTS_TRUNCATE (
    contents : Contents_handle
)

```

CONTENTS\_TRUNCATE truncates *contents* from the current position to the end.

The "contents\_size" attribute of *contents* is set to indicate the new size.

LAST is reset to one less than the current position, which is unchanged, except when the current position is FIRST, in which case LAST is undefined and the file is empty.

This operation applies only to files.

### Errors

CONFIDENTIALITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED (object determined by *contents*, ATOMIC)

CONFIDENTIALITY\_WOULD\_BE\_VIOLATED (object determined by *contents*, ATOMIC)

CONTENTS\_IS\_NOT\_FILE\_CONTENTS (*contents*)

CONTENTS\_IS\_NOT\_OPEN (*contents*)

CONTENTS\_OPERATION\_IS\_INVALID (*contents*)

INTEGRITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED (object determined by *contents*, ATOMIC)

INTEGRITY\_WOULD\_BE\_VIOLATED (object determined by *contents*, ATOMIC)

OBJECT\_IS\_INACCESSIBLE (object determined by *contents*, ATOMIC)

VOLUME\_IS\_FULL (volume containing object determined by *contents*)

VOLUME\_IS\_INACCESSIBLE (volume containing object determined by *contents*, ATOMIC)

NOTE - CONTENTS\_TRUNCATE can affect the size of a file while other operations are accessing it.

## 12.2.12 CONTENTS\_WRITE

```

CONTENTS_WRITE (
    contents      : Contents_handle,
    data          : Unstructured_contents
)
    actual_size   : Natural
  
```

CONTENTS\_WRITE writes some or all of a sequence of octets *data* to *contents* at the current position, and returns the number of octets actually written.

If *contents* is a file with opening mode READ\_WRITE, WRITE\_ONLY, or APPEND\_ONLY and if writing to the file would not cause its size to exceed the MAX\_FILE\_SIZE, the sequence of octets *data* is written from the current position, and the current position is changed to the position following the last written octet. The contents size of *contents* is set to indicate the new size.

If *contents* is a pipe with opening mode APPEND\_ONLY and if writing to the pipe would not cause its size to exceed MAX\_PIPE\_SIZE, the sequence of octets *data* is written from the position LAST + 1, and the current position is changed to the position following the last written octet.

If *contents* is a device with opening mode READ\_WRITE, WRITE\_ONLY, or APPEND\_ONLY and if writing to the device would not cause its size to exceed any device-dependent maximum size limit, the sequence of octets *data* is written from the current position, and the current position is changed to the position following the last written octet.

If the available space does not allow the whole of *data* to be written to *contents*:

- if *contents* is a file and at least one octet can be written, as many octets as possible are written;
- if *contents* is a file and no octet can be written (e.g. MAX\_FILE\_SIZE has been reached), the operation fails;
- if *contents* is a non-blocking pipe, as many octets from *data* as there is room for are written; if no octets can be written (i.e. MAX\_PIPE\_SIZE has been reached) the operation fails;
- if *contents* is a pipe which is not non-blocking, the operation waits, but on normal completion (i.e. after space has been made available in the pipe and no interrupt occurred) the operation has written all the octets;
- if *contents* is a device which is not non-blocking, the operation waits until octets can be written;
- if *contents* is a non-blocking device, as many octets as there are room for are written; if no octet can be written, the operation fails.

In all cases, the octets of *data* are written in order starting with the first element, and the actual number of octets written to *contents* is returned in *actual\_size*.

If a concurrent CONTENTS\_TRUNCATE operation is performed on the object *contents* after *contents* is opened, *data* is nevertheless written at the specified current position (i.e. there is no interference implying that the current position is reset) as if prior to the write a CONTENTS\_SEEK operation had been performed with the current position as argument.

### Errors

CONFIDENTIALITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED (object determined by *contents*, ATOMIC)

CONFIDENTIALITY\_WOULD\_BE\_VIOLATED (object determined by *contents*, ATOMIC)

CONTENTS\_IS\_NOT\_OPEN (*contents*)

CONTENTS\_OPERATION\_IS\_INVALID (*contents*)

DEVICE\_LIMIT\_WOULD\_BE\_EXCEEDED (*data*, *contents*)

INTEGRITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED (object determined by *contents*, ATOMIC)

INTEGRITY\_WOULD\_BE\_VIOLATED (object determined by *contents*, ATOMIC)

LIMIT\_WOULD\_BE\_EXCEEDED ((MAX\_FILE\_SIZE, MAX\_PIPE\_SIZE))

OBJECT\_IS\_INACCESSIBLE (object determined by *contents*, ATOMIC)

If *contents* is a file:

PROCESS\_FILE\_SIZE\_LIMIT\_WOULD\_BE\_EXCEEDED (*data*, *contents*)

VOLUME\_IS\_FULL (volume containing object determined by *contents*)

VOLUME\_IS\_INACCESSIBLE (volume containing object determined by *contents*)

### 12.2.13 DEVICE\_GET\_CONTROL

```

DEVICE_GET_CONTROL (
    contents      : Contents_handle,
    operation     : Natural
)
    control_data : Control_data

```

DEVICE\_GET\_CONTROL returns control information from the device contents *contents* in *control\_data*, according to *operation*. The meanings of *operation* and *control\_data* are implementation-defined and may be device-dependent.

#### Errors

CONFIDENTIALITY\_WOULD\_BE\_VIOLATED (object determined by *contents*, ATOMIC)

CONTENTS\_IS\_NOT\_OPEN (*contents*)

DEVICE\_CONTROL\_OPERATION\_IS\_INVALID (*contents*, *operation*)

INTEGRITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED (object determined by *contents*, ATOMIC)

OBJECT\_IS\_INACCESSIBLE (volume containing object determined by *contents*, ATOMIC)

VOLUME\_IS\_INACCESSIBLE (volume on which object determined by *contents* resides, ATOMIC)

### 12.2.14 DEVICE\_SET\_CONTROL

```

DEVICE_SET_CONTROL (
    contents      : Contents_handle,
    operation     : Natural,
    control_data  : Control_data
)

```

DEVICE\_SET\_CONTROL performs a control operation on the device contents *contents*, according to *operation*. The parameters for the operation are specified in *control\_data*. The meanings of *operation* and *control\_data* are implementation-defined and may be device-dependent.

#### Errors

CONFIDENTIALITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED (object determined by *contents*, ATOMIC)

CONFIDENTIALITY\_WOULD\_BE\_VIOLATED (object determined by *contents*, ATOMIC)

CONTENTS\_IS\_NOT\_OPEN (*contents*)

DEVICE\_CONTROL\_OPERATION\_IS\_INVALID (*contents*, *operation*)

INTEGRITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED (object determined by *contents*, ATOMIC)

INTEGRITY\_WOULD\_BE\_VIOLATED (object determined by *contents*, ATOMIC)

OBJECT\_IS\_INACCESSIBLE (object determined by *contents*, ATOMIC)

VOLUME\_IS\_FULL (volume containing object determined by *contents*)

VOLUME\_IS\_INACCESSIBLE (volume containing object determined by *contents*, ATOMIC)

## 13 Process execution

### 13.1 Process execution concepts

#### 13.1.1 Static contexts

**sds** system:

static\_context: **child type of file with attribute**

max\_inheritable\_open\_objects: **natural** := 3;

interpretable: **boolean** := **false**;

**link**

interpreter: **reference link to** static\_context;

restricted\_execution\_class: **reference link to** execution\_class;

**end** static\_context;

**end** system;

The max (maximum number of) inheritable open objects is the maximum number of open objects that a process running the static context may inherit from the process which created it.

A static context is *interpretable* if "interpretable" is **true**; otherwise it is *executable*.

The *interpreter* of an interpretable static context is the destination of the "interpreter" link, if there is one; it must not itself be interpretable.

A static context is *foreign* if it has a restricted execution class and that execution class (see 13.1.3) has a usable execution site which is a foreign system; otherwise it is *native*.

The *execution class* of an executable static context is the set of execution sites in which the static context may run. If the static context has a "restricted\_execution\_class" link then its execution class contains just the destination object of that link; otherwise it contains all the execution sites in the PCTE installation. The execution class of an interpretable static context is the intersection of that set and the execution class of the actual interpreter of the static context.

#### NOTES

1 A static context (short for static context of a program) is an executable or interpretable program in a static form that can be run by a process, either directly by loading and executing it (executable) or indirectly by running another static context as an interpreter (interpretable). It may be run either by a PCTE implementation or by a foreign system.

2 The default of 3 for maximum inheritable open objects allows inheritance of standard input, output and error channels as supported by some operating systems. The number of open objects is limited to MAX\_OPEN\_OBJECTS\_PER\_PROCESS (see clause 24) so this is the maximum effective value for maximum inheritable open objects.

3 The format of the contents of an executable static context is implementation-defined by the PCTE implementation (for workstations in the execution class) or the foreign system implementation (for foreign systems in the execution class) of the execution site.

4 If an interpretable static context has no interpreter, a static context is selected to interpret it as described in PROCESS\_START.

5 A static context has other properties defined in the security SDS (see 19.1.1).

6 The fact that the interpreter of an interpretable static context is not interpretable is checked by PROCESS\_START and PROCESS\_CREATE\_AND\_START, but not by LINK\_CREATE, OBJECT\_SET\_ATTRIBUTE, OBJECT\_SET\_SEVERAL\_ATTRIBUTES, OBJECT\_RESET\_ATTRIBUTE, or OBJECT\_DELETE\_ATTRIBUTE.



### 13.1.2 Foreign execution images

**sds** system:

foreign\_execution\_image: **child type of object with attribute**

foreign\_name: **string**;

**link**

on\_foreign\_system: **reference link to foreign\_system**;

**end** foreign\_execution\_image;

**end** system;

The syntax and semantics of the foreign name are implementation-defined.

The "on\_foreign\_system" link defines a foreign system which may execute the foreign execution image.

#### NOTES

- 1 A foreign execution image differs from a static context for use on a foreign system in that it is only a representation of the image to be executed. The execution image itself is of undefined format and is not represented in the object base.
- 2 The foreign name is intended to provide enough information to determine the foreign system object, e.g. a file, which contains an execution image.

### 13.1.3 Execution classes

**sds** system:

execution\_site\_identifier: **natural**;

execution\_class: **child type of object with link**

usable\_execution\_site: **reference link** (execution\_site\_identifier) **to** execution\_site;

**end** execution\_class;

**end** system;

An execution class specifies a set of execution sites (workstations or foreign systems) on which any static context with that execution class may be executed. Execution sites are defined in 18.1.

#### NOTES

- 1 If a static context has no restricted execution class, the choice of execution site may be specified when the static context is run; otherwise it is implementation-defined.
- 2 If an execution class has no usable execution site, a static context with that execution class as a restricted execution class is unable to run. Thus it is possible to (temporarily) prevent a static context from running.
- 3 The addition and removal of execution sites to and from an execution class is performed using operations of clause 9. An execution site is a usable execution site of an execution class if and only if there is a "usable\_execution\_site" link between the site and the class. The value of the key of such a link is unimportant.
- 4 While it is recommended that tools keep the "execution\_site\_identifier" key consistent with the execution site identifier of the usable execution site in the execution site directory, a PCTE implementation is not required to enforce this consistency, nor even to ensure that the key is any execution site identifier in the execution site directory.
- 5 The definition of an execution class allows both workstations and foreign systems to be of the same execution class. In practice, such a mixed class is unlikely to be useful.

### 13.1.4 Processes

Initial\_status = RUNNING | SUSPENDED | STOPPED

**sds** system:

inheritable: **boolean** := **true**;

referenced\_object: (**navigate**) **designation link** (reference\_name: **string**) to object with **attribute**

inheritable;

**end** referenced\_object;

open\_object: (**navigate**) **designation link** (open\_object\_key: **natural**) to file, pipe, device with **attribute**

opening\_mode: (**read**) **enumeration** (READ\_WRITE, READ\_ONLY, WRITE\_ONLY, APPEND\_ONLY) := READ\_ONLY;

non\_blocking\_io: (**read**) **boolean**;

inheritable;

**end** open\_object;

is\_listener: (**navigate**) **non\_duplicated designation link** (number) to

message\_queue with **attribute**

message\_types: (**read**) **string**;

**end** is\_listener;

process\_waiting\_for: (**navigate**) **designation link** (number) to object with **attribute**

waiting\_type: (**read**) **enumeration** (WAITING\_FOR\_LOCK, WAITING\_FOR\_TERMINATION, WAITING\_FOR\_WRITE, WAITING\_FOR\_READ) := WAITING\_FOR\_LOCK;

locked\_link\_name;

**end** process\_waiting\_for;

process: **child type of object with attribute**

process\_status: (read) **non\_duplicated enumeration** (UNKNOWN, READY, RUNNING, STOPPED, SUSPENDED, TERMINATED) := UNKNOWN;  
 process\_creation\_time: (read) **time**;  
 process\_start\_time: (read) **time**;  
 process\_termination\_time: (read) **time**;  
 process\_user\_defined\_result: **string**;  
 process\_termination\_status: (read) **integer**;  
 process\_priority: (read) **natural**;  
 process\_file\_size\_limit: (read) **natural**;  
 process\_string\_arguments: (read) **string**;  
 process\_environment: (read) **string**;  
 process\_time\_out: (read) **natural**;  
 acknowledged\_termination: (read) **boolean**;  
 deletion\_upon\_termination: (read) **boolean** := true;  
 time\_left\_until\_alarm: (read) **non\_duplicated natural**;  
 character\_encoding: (read) **non\_duplicated natural**;

**link**

process\_object\_argument: **designation link** (number) **to** object;  
 executed\_on: (navigate) **designation link to** execution\_site;  
 referenced\_object;  
 open\_object;  
 reserved\_message\_queue: (navigate) **designation link** (number) **to** message\_queue;  
 is\_listener;  
 default\_interpreter: **designation link to** static\_context;  
 actual\_interpreter: (navigate) **designation link to** static\_context;  
 process\_waiting\_for;  
 parent\_process: (navigate, delete) **implicit link to** process **reverse** child\_process;  
 started\_in\_activity: (navigate) **reference link to** activity **reverse** process\_started\_in;

**component**

child\_process: (navigate, delete) **composition link** (number) **to** process **reverse** parent\_process;  
 started\_activity: (navigate) **composition link** (number) **to** activity **reverse** started\_by;

**end process;**

**end system;**

**sds metasds:**

**import object type** system-process;

**extend object type** process **with**

**link**

sds\_in\_working\_schema: (navigate) **designation link** (number) **to** sds;

**end process;**

**end metasds;**

A process is a means of running a static context or foreign execution image. *Creation* of a process refers to the action of PROCESS\_CREATE. A process *runs* the static context (executable or interpretable) or foreign execution image specified when the process is created. A process *executes* the static context or foreign execution image specified when the process is created unless an interpretable static context is specified, in which case it executes another static context which is executable.

A process executes by the execution of one or more *threads*. Within a process, threads may execute in parallel (proceed independently), or execution may switch between threads, or both,

according to rules not defined in this part of ISO/IEC 13719. A thread is *suspended* (i.e. its execution does not progress) when it is executing an operation which is waiting for the occurrence of an event (see 8.7.2). A binding must define the mapping of threads and of their suspension to the binding language. A binding may impose limitations on threads by the definition of the rules for their interaction; in particular, a binding may specify that, except for the activation or waking of a handler (see 14.1), a process always executes by the execution of one and only one thread. The activation of a handler normally involves execution of a separate thread, although there may be special binding-defined rules governing this execution.

Whether other threads of a process (if any) can continue to execute while one thread is suspended, and whether such threads can issue from operation calls, are instances of binding-defined rules governing the execution of threads.

The process status is the status of the process with respect to execution. State transitions occur as the result of operations in 13.2 or of events outside tool control, e.g. a thread reaching a breakpoint. The process status may have the following values:

- READY: ready to execute.
- RUNNING: executing: one or more threads of the process are running or suspended.
- STOPPED: stopped from execution: all threads of the process are stopped; this is for use in process monitoring, see 13.5.
- SUSPENDED: suspended from execution: all threads of the process are suspended; this results from PROCESS\_SUSPEND.
- TERMINATED: prevented from further execution.

The status value STOPPED is required only by the monitoring module (see 13.5).

In addition the process status has an initial value UNKNOWN which it is given if the process is created by operations in clause 9. Such a process is prevented from executing.

If one thread of a process is stopped, then all are, and similarly with suspension.

The terms *ready*, *running*, *stopped*, *suspended*, *terminated* and *unknown* apply to a process whose process status is READY, RUNNING, STOPPED, SUSPENDED, TERMINATED or UNKNOWN, respectively. The terms 'running' and 'stopped' are also applied to a thread of a process. A process *starts* when its status changes from READY.

A *breakpoint* is an implementation-defined marker defining a point in a process such that when execution reaches that point while the process is running, the process status is changed to STOPPED.

If a handler is executed on behalf of a suspended listening process, the status of the listening process is changed to RUNNING (see 14.1).

The precise time of a change of process status as recorded in the process creation, start or termination time is implementation-dependent except that it is between the start and end of any operation that causes the change of process status.

The process creation time is the time when the process was created.

The process start time is the time when the process started to run a static context or foreign execution image. Its value is the default value of time attributes if the process is ready.

The process termination time is the time when the process terminated. Its value is the default value of time attributes unless the process is terminated.

The semantics of the process user defined result are not defined in this part of ISO/IEC 13719.

The process termination status specifies the conditions under which the process terminated. Its value is the default value of integer attributes unless the process is terminated. The process termination status has two sets of named values, whose actual values are implementation-defined. Other values may be set using `PROCESS_SET_TERMINATION_STATUS` or `PROCESS_TERMINATE` but are not defined in this Standard. The sets of named values are:

- Success:
  - `EXIT_SUCCESS`: The process has terminated normally, i.e. not as in the failure cases. The process termination status has this value if a process terminates other than by `PROCESS_TERMINATE` (explicitly or implicitly called), and the process termination status has not been changed by `PROCESS_SET_TERMINATION_STATUS`.
- Failure:
  - `EXIT_ERROR`: The process has been terminated abnormally by itself (using `PROCESS_TERMINATE`).
  - `FORCED_TERMINATION`: The process has been terminated abnormally by another process (using `PROCESS_TERMINATE`).
  - `SYSTEM_FAILURE`: The process has been terminated abnormally by the PCTE implementation.
  - `ACTIVITY_ABORTED`: The process has been terminated abnormally as a result of the destination of its "started\_in\_activity" link being aborted by `ACTIVITY_ABORT`.

The process priority defines the priority of running the process relative to that of other processes. The range of values is from 0 to the implementation-defined limit `MAX_PRIORITY_VALUE`. Their effect is implementation-defined except that a greater integer value indicates a greater priority.

The process file size limit defines the maximum contents size of each file to which the process writes.

The value of the process string arguments is a string of the following syntax, which defines it as a sequence of zero or more substrings. Each substring is an argument preceded by the length of the argument in hexadecimal notation. The semantics of the sequence of arguments is not defined in this part of ISO/IEC 13719.

```
arguments = {substring};
substring= length, argument;
length = hex digit, hex digit, hex digit, hex digit;
hex digit = '0' | '1' | '2' | '3' | '4' | '5' | '6' | '7' | '8' | '9' | 'A' | 'B' | 'C' | 'D' | 'E' | 'F';
argument = (*any sequence of graphic characters*);
```

The semantics of process environment is not defined in this part of ISO/IEC 13719. The value has the same syntax as the process string arguments.

The process time out limits the duration of each *indivisible* operation, i.e. each operation whose execution does not cause it to wait, and also of each operation whose execution causes the creation of a "process\_waiting\_for" link from the calling process. If the value is 0, the limit is infinite, otherwise the limit is the value in seconds. An operation whose duration exceeds the limit terminates with the error OPERATION\_HAS\_TIMED\_OUT.

If the value is greater than 0, the time left until alarm defines the maximum duration in seconds that a process will be suspended when it next suspends or, while the process is suspended, the maximum duration until it is resumed. If the process is resumed before the alarm goes off, the value of time left until alarm indicates the unexpired duration. Otherwise when the time expires the process receives an implementation-defined alarm message of message type WAKE (provided it has reserved a message queue and is handling wakeup messages) and is resumed. If there is more than one reserved message queue that has a handler enabled to handle WAKE messages, the system chooses in an implementation-defined way which message queue receives the WAKE message.

The acknowledged termination is **true** when the process has terminated and the parent process has continued running after waiting for termination.

If deletion upon termination is **true** and the deletion conditions are satisfied, the process is deleted automatically when it terminates, after acknowledged termination of this process has been set **true** by the parent process.

The character encoding defines the single- or multi-octet character set in terms of which pathnames, keys, and type names are interpreted. The default value of 0 refers to ISO 8859-1; the interpretation of other values is implementation-defined. An unrecognised value is defaulted to 0.

The "sds\_in\_working\_schema" links specify by their key values a sequence of SDSs which determines the working schema of the process (see 8.1). The "sds\_in\_working\_schema" links are created when a process is created and may be changed by PROCESS\_SET\_WORKING\_SCHEMA.

The semantics of the process object arguments is not defined in this part of ISO/IEC 13719.

The destination of the "executed\_on" link is called the *execution site of the process*.

Referenced objects are used in the construction of object designators through their reference names (see clause 23). Referenced objects are created and deleted by PROCESS\_SET\_REFERENCED\_OBJECT and PROCESS\_UNSET\_REFERENCED\_OBJECT respectively. Reference name values are restricted to the values of key string value defined in 23.1.2.7.

The following reference names are reserved and refer to the given referenced objects:

- 'self': This process. This referenced object always exists, cannot be changed and has inheritability **false**.
- 'static\_context': The static context run by the process. This referenced object always exists, cannot be changed and has inheritability **false**.
- 'common\_root': The common root (see 9.1.2). This referenced object always exists, cannot be changed and has inheritability **true**.
- 'home\_object': The meaning of the 'home\_object' referenced object is not defined in this part of ISO/IEC 13719.

- 'current\_object': The meaning of the 'current\_object' referenced object is not defined in this part of ISO/IEC 13719. Conventions for using it are given in 23.1.2.2.

The referenced objects with reference names "static\_context", "common\_root", "home\_object", and "current\_object" are known as the *static context of the process*, *common root*, *home object*, and *current object* respectively.

If inheritability is **true** the referenced object is to be made a referenced object of each child process created by this process (and inheritability is to be set **true** for it). The inheritability of a referenced object may be changed by operations in clause 9. An inherited "referenced\_object" link may be deleted by the child process but this does not affect the referenced objects of the creating process.

An open object is an object opened for access to its contents (see clause 12). If inheritable is set **true**, the open object is to be inherited as opened (and the corresponding current position is to be shared) by each child process created by this process in the manner specified by the attributes of the "open\_object" link (and inheritable is to be set **true** for the child's open object). The inheritable attribute of an open object may be changed by operations in clause 9. An inherited open object may be closed by the child process but this does not affect the open objects of the creating process. The semantics of the other attributes of an open object are defined by the operations in clause 12.

For open objects with keys 0, 1 and 2 see 12.1.

The default interpreter, if it exists, is a static context which will interpret the static context run by a process if it is interpretable and has no interpreter. The value of the default interpreter may be changed by operations in clause 9.

The actual interpreter is the static context that interprets an interpretable static context.

For reserved message queue and "is\_listener" see 14.1.

The destination of the "process\_waiting\_for" link is a resource that a thread of the process is waiting for; for further attributes see 16.1.2. A link of this type exists for each operation that is waiting. The waiting type values are:

- WAITING\_FOR\_LOCK: waiting to establish a lock on a resource which already has an incompatible lock.
- WAITING\_FOR\_TERMINATION: waiting for a child process to terminate.
- WAITING\_FOR\_WRITE: waiting to write to a full message queue, a full pipe, a device, an audit file or an accounting log.
- WAITING\_FOR\_READ: waiting to read from a message queue containing no message of the specified type, an empty pipe, or a device.

The started in activity is the activity which was the current activity of the parent process at the time the process was created. Activities are defined in clause 16.

A process is either the initial process of a workstation (see 18.1.2) or a child process of one other process.

The parent process is the process which created this process or another process nominated by the creating process to be the parent.

For the started activities, see clause 16.



## NOTES

- 1 The process user defined result is provided for tool-defined use, especially for a child process to pass back results to its parent on termination.
- 2 The process priority is intended to be mapped to the process priority of an underlying operating system (if there is one). The number of possible values should be a power of 2.
- 3 The process string arguments is intended for passing parameters in the form of strings to a child process running a tool written in a language which specifies a mechanism for passing parameters to the tool. The specification of the length in hexadecimal notation enables the maximum length of an argument to be stored in 2 octets.
- 4 The process environment is provided as a mechanism for modifying aspects of the environment in which a child process is to run.
- 5 If the acknowledged termination of a process is **true**, then the process has terminated but could not be deleted, e.g. because deletion upon termination is **false**, or because there was a reference link to the process.
- 6 The "process\_object\_argument" link is intended for designating an object to a process, e.g. a print spooler, while it is running. The process may use key values to distinguish the different objects so designated if it does not delete the link to each process object argument after it has been processed.
- 7 A process can only be moved (thus changing its volume number) while the process status is READY or TERMINATED. It is recommended that a ready process is only moved to a volume that is controlled by the execution site which is to execute the process or, if the execution site is a discless workstation, to one that can be accessed efficiently.
- 8 The child processes of a process are components of that process, but this does not mean that operations on a process apply also to its child processes; e.g. terminating a process does not of itself terminate its child processes.
- 9 Many of the links of process that have no reverse link have a corresponding link which is effectively a reverse link except that only the link from the process exists before a process is started.
- 10 Operations specific to processes, i.e. those with names starting with "PROCESS\_", do not establish any locks on the process, its links or its attributes (and thus these changes are not reversed if the transaction is aborted).
- 11 Operations specific to processes do not require discretionary access control on the calling process, its links or its attributes.
- 12 A process has other properties defined in the security and accounting SDSs.
- 13 The implicit creation and deletion of a usage designation link is allowed by operations defined in clause 13 even if the origin object of the link resides on a read-only volume or is a copy object.
- 14 Table 3 shows the available transitions of process status.

Table 3 - Available transitions of process status

| From                                                                                                                                                                                                                                                                                                                                                | To    |         |           |         |            |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------|---------|-----------|---------|------------|
|                                                                                                                                                                                                                                                                                                                                                     | Ready | Running | Suspended | Stopped | Terminated |
| (nonexistent)                                                                                                                                                                                                                                                                                                                                       | CR    | CS      | X         | X       | X          |
| Ready                                                                                                                                                                                                                                                                                                                                               | N     | ST      | ST        | ST      | X          |
| Running                                                                                                                                                                                                                                                                                                                                             | X     | N       | SU        | BP      | TE         |
| Suspended                                                                                                                                                                                                                                                                                                                                           | X     | RE,H    | N         | X       | TE         |
| Stopped                                                                                                                                                                                                                                                                                                                                             | X     | CO      | X         | N       | TE         |
| Terminated                                                                                                                                                                                                                                                                                                                                          | X     | X       | X         | X       | N          |
| <b>Key</b><br>BP      breakpoint<br>CO      PROCESS_CONTINUE<br>CR      PROCESS_CREATE<br>CS      PROCESS_CREATE_AND_START<br>RE      PROCESS_RESUME<br>ST      PROCESS_START<br>SU      PROCESS_SUSPEND<br>TE      PROCESS_TERMINATE<br>H      execution of a message handler (see 14.1)<br>N      null transition<br>X      impossible transition |       |         |           |         |            |

15 It is intended that a PCTE implementation maintain its integrity against operation calls from concurrent threads. In addition, the implementation may provide some degree of concurrency within operations, but that is not mandatory. Thus operations called in concurrent threads may block immediately until an operation called earlier has terminated. Such implementation dependence is likely to apply to all language bindings supported by the implementation in addition to binding dependences that result from the level of support for threads by the binding language.

16 When a new process is created, it inherits the value of the character set of the calling process. Thus an implementation can set the character set of the initial process to the preferred value; as all processes are derived from the initial process, they would then work by default with this character set, and no further action would be necessary to select the character set, unless another character set is required for some tools.

### 13.1.5 Initial processes

Each workstation in a PCTE installation has an *initial process*; this is a process that is created by implementation-dependent means such that, when it starts to run a tool, it is indistinguishable from a process that has been created by PROCESS\_CREATE and modified by other PCTE operations, except that the initial process has no parent process. When the first static context runs in the initial process, the initial process has the following particular values for attributes and links:

- the volume on which the process resides is the administration volume of the execution site of the initial process;

- the execution site of the process is the workstation for which the process is the initial process;
- the static context of the process is the static context being run by the initial process;
- the destination of the "actual\_interpreter" link is the static context being executed by the initial process, if any;
- the destination of the "started\_in\_activity" is the outermost activity of the execution site (see 16.1.1);
- the static context of the initial process is a member of the predefined program group PCTE\_SECURITY or of a program group which has PCTE\_SECURITY as one of its program supergroups.

NOTE - The initial process of a workstation is intended to start one or more processes, each of which runs a static context, typically a login or user authentication tool (which may be a portable tool), to perform various tasks when a human user starts or ends a session at the workstation. It has no consumer identity. The tasks to be performed at the start of the session may include, for example:

- authenticating the human user and setting the discretionary and mandatory context appropriate to that user by calling PROCESS\_SET\_USER\_AND\_USER\_GROUP\_IDENTITY; this must be done before any processing on behalf of the user to assure the security of the PCTE installation;
- initializing a general purpose environment for the running of tools by the user;
- tailoring the environment to the user, for example by setting the referenced object "home\_object".

### 13.1.6 Profiling and monitoring concepts

Profile\_handle :: Token

Buffer = seq of Natural

Address :: Token

Process\_data = seq of Octet

These types are used in profiling and monitoring operations; see 13.4 and 13.5.

## 13.2 Process execution operations

### 13.2.1 PROCESS\_CREATE

```

PROCESS_CREATE (
  static_context      : Static_context_designator | Foreign_execution_image_designator,
  process_type        : Process_type_nominator,
  parent              : [ Process_designator ],
  site                : [ Execution_site_designator ],
  implicit_deletion    : Boolean,
  access_mask         : Atomic_access_rights
)
new_process          : Process_designator

```

If no value is supplied for *parent*, *parent* designates the calling process.

PROCESS\_CREATE creates a process that is able to run a static context. The new process becomes a child process of *parent* (either the calling process or an ancestor of the calling process).

The new process *new\_process* is of type *process\_type* with attributes and links as defined below.

- Attributes and links of type "object" defined in SDS 'system' as by OBJECT\_CREATE, except that "volume\_identifier" is set to "volume\_identifier" of *parent*, if *parent* and the new process have the same execution site, otherwise to "volume\_identifier" of the new process's execution site.
- Attributes and links of type "process" defined in SDS 'system':
  - . "process\_status" is set to READY;
  - . "process\_creation\_time" is set to the current time (a value of system time between the start and end of the operation);
  - . "process\_priority" is set to "process\_priority" of the calling process;
  - . "process\_file\_size\_limit" is set to "process\_file\_size\_limit" of the calling process;
  - . "deletion\_upon\_termination" is set to *implicit\_deletion*;
  - . the character encoding is set to the character encoding of the calling process;
  - . "sds\_in\_working\_schema" links are created, each with the same destination and key as each of the "sds\_in\_working\_schema" links of the calling process;
  - . an "executed\_on" link is created to *site*, or if *site* is absent:
    - . if *static\_context* is executable, to an implementation-dependent member of the execution class of *static\_context*;
    - . if *static\_context* is interpretable, to an implementation-dependent member of the intersection of the execution classes of *static\_context* and its interpreter;
    - . if *static\_context* is a foreign execution image, to the destination of the "on\_foreign\_system" link from *static\_context*;
  - . for each "referenced\_object" link of the calling process with inheritability **true** (except for the referenced objects "self" and "static\_context") a "referenced\_object" link is created with the same destination and key; in addition, "referenced\_object" links with reference names "self" and "static\_context" are created with destinations the new process and *static\_context*, respectively, and inheritability **false**;
  - . for each "open\_object" link of the calling process with inheritable **true** an "open\_object" link is created with the same destination and key, and with the same opening mode and non-blocking io, in ascending order of key value, up to a limit of "max\_inheritable\_open\_objects" of *static\_context*;
  - . if the calling process has a default interpreter, a "default\_interpreter" link is created to the default interpreter of the calling process;
  - . a "parent\_process" link to *parent* and its reverse "child\_process" link are created;
  - . a "started\_in\_activity" link to the current activity of *parent* and its reverse "process\_started\_in" link are created.
- Attributes and links of type "object" defined in SDS 'discretionary\_security' as by OBJECT\_CREATE with *access\_mask*, except:

- "atomic\_acl" has two additional groups added, if not already present, and in any case these two groups are granted all access rights. These groups are:
    - the user of the new process;
    - the predefined security group PCTE\_EXECUTION;
  - "confidentiality\_label" is set to "confidentiality\_label" of the calling process;
  - "integrity\_label" is set to "integrity\_label" of the calling process.
- Attributes and links of type "process" defined in SDSs 'discretionary\_security' and 'mandatory\_security':
  - "default\_atomic\_acl" is set to "default\_atomic\_acl" of the calling process;
  - "default\_object\_owner" is set to "default\_object\_owner" of the calling process;
  - "floating\_confidentiality\_level" is set to "floating\_confidentiality\_level" of the calling process;
  - "floating\_integrity\_level" is set to "floating\_integrity\_level" of the calling process;
  - a "user\_identity" link is created to the user of the calling process;
  - an "adopted\_user\_group" link is created to the adopted user group of the calling process;
  - "adoptable\_user\_group" links are created, each with the same destination and key as each of those "adoptable\_user\_group" links of the calling process with "adoptable\_for\_child" **true**.
- Attributes and links of type "process" defined in SDS 'accounting':
  - a "consumer\_identity" link is created to the consumer identity of the calling process, if any.

PROCESS\_CREATE returns a designator of the new process as *new\_process*.

If the workstation controlling the device on which is mounted the volume on which *new\_process* resides becomes inaccessible before *new\_process* is started, the "sds\_in\_working\_schema", "executed\_on", "opened\_objects", "user\_identity", "adopted\_user\_group", "adoptable\_user\_group", "referenced\_object", and "consumer\_identity" designation links from *new\_process* are deleted, the status of *new\_process* is set to TERMINATED and the exit status of *new\_process* is set to SYSTEM\_FAILURE.

### Errors

ACCESS\_ERRORS (*static\_context*, ATOMIC, MODIFY, EXECUTE)

If *static\_context* is interpretable:

ACCESS\_ERRORS (interpreter of *static\_context*, ATOMIC, READ, EXECUTE)

ACCESS\_ERRORS (*parent*, ATOMIC, READ, APPEND\_LINKS)

ACCESS\_ERRORS (the current activity of *parent*, ATOMIC, MODIFY, APPEND\_IMPLICIT)

EXECUTION\_CLASS\_HAS\_NO\_USABLE\_EXECUTION\_SITES (execution class of *static\_context*)

EXECUTION\_SITE\_IS\_INACCESSIBLE (*site*)

EXECUTION\_SITE\_IS\_NOT\_IN\_EXECUTION\_CLASS (*site*, *static\_context*)

EXECUTION\_SITE\_IS\_UNKNOWN (*site*)

If *static\_context* is a foreign execution image:

FOREIGN\_EXECUTION\_IMAGE\_HAS\_NO\_SITE (*static\_context*)

LABEL\_IS\_OUTSIDE\_RANGE (*new\_process*, the volume on which *new\_process* would reside)

LABEL\_IS\_OUTSIDE\_RANGE (*new\_process*, the would-be execution site of *new\_process*)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_PROCESSES)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_PROCESSES\_PER\_USER)

OBJECT\_OWNER\_VALUE\_WOULD\_BE\_INCONSISTENT\_WITH\_ATOMIC\_ACL

If *parent* has OWNER granted or denied:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION (*new\_process*)

PROCESS\_LACKS\_REQUIRED\_STATUS (*parent*, (READY, RUNNING, STOPPED, SUSPENDED))

If *parent* is not the calling process:

PROCESS\_IS\_NOT\_ANCESTOR (*parent*)

PROCESS\_IS\_UNKNOWN (*parent*)

REFERENCE\_CANNOT\_BE\_ALLOCATED

STATIC\_CONTEXT\_REQUIRES\_TOO\_MUCH\_MEMORY (*static\_context*)

USAGE\_MODE\_ON\_OBJECT\_TYPE\_WOULD\_BE\_VIOLATED ("object", *process\_type*)

If *process* is the calling process:

VOLUME\_IS\_FULL (calling process)

NOTE - It is implementation-dependent which underlying resources (e.g. memory) required for process execution are allocated by PROCESS\_CREATE and which are allocated by PROCESS\_START.

### 13.2.2 PROCESS\_CREATE\_AND\_START

```
PROCESS_CREATE_AND_START (
    static_context      : Static_context_designator | Foreign_execution_image_designator,
    arguments           : String,
    environment        : String,
    site                : [ Execution_site_designator ],
    implicit_deletion   : Boolean,
    access_mask         : Atomic_access_rights
)
new_process           : Process_designator
```

PROCESS\_CREATE\_AND\_START creates and runs a process asynchronously in one single operation.

The overall effect is as for the following sequence of operations.

```
new_process := PROCESS_CREATE (static_context, "process", nil, site,
    implicit_deletion, access_mask);
```

```
PROCESS_START (new_process, arguments, environment, site, RUNNING);
```

PROCESS\_CREATE\_AND\_START is an atomic operation for the calling process.

**Errors**

ACCESS\_ERRORS (*static\_context*, ATOMIC, MODIFY, EXECUTE)

For each SDS *sds* which is the destination of an "in\_working\_schema\_of" link from *new\_process*:

ACCESS\_ERRORS (*sds*, ATOMIC, SYSTEM\_ACCESS)  
 DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED\_TO\_PROCESS  
 (*new\_process*, *sds*, ATOMIC, EXPLOIT\_SCHEMA)

For each open object *object* which is the destination of an "open\_object" link from the calling process:

ACCESS\_ERRORS (*object*, ATOMIC, SYSTEM\_ACCESS)  
 If the link's opening mode attribute is READ\_ONLY or READ\_WRITE,  
 DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED\_TO\_PROCESS  
 (*new\_process*, *object*, ATOMIC, READ, READ\_CONTENTS)  
 If the link's opening mode attribute is WRITE\_ONLY or READ\_WRITE,  
 DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED\_TO\_PROCESS  
 (*new\_process*, *object*, ATOMIC, READ, WRITE\_CONTENTS)  
 If the link's opening mode attribute is APPEND\_ONLY,  
 DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED\_TO\_PROCESS  
 (*new\_process*, *object*, ATOMIC, READ, APPEND\_CONTENTS)

For the activity *activity* which is the destination of the "started\_in\_activity" link from the calling process:

ACCESS\_ERRORS (*activity*, ATOMIC, SYSTEM\_ACCESS)  
 ACTIVITY\_STATUS\_IS\_INVALID (*activity*, ACTIVE)

For the user *user* which is the destination of the "user\_identity" link from the calling process:

ACCESS\_ERRORS (*user*, ATOMIC, SYSTEM\_ACCESS)

For the user group *group* which is the destination of the "adopted\_user\_group" link from the calling process:

ACCESS\_ERRORS (*group*, ATOMIC, SYSTEM\_ACCESS)

For the consumer group *group* which is the destination of the "consumer\_identity" link from the calling process:

ACCESS\_ERRORS (*group*, ATOMIC, SYSTEM\_ACCESS)

If *static\_context* has an interpreter:

ACCESS\_ERRORS (interpreter of *static\_context*, ATOMIC, MODIFY, EXECUTE)

CONTROL\_WOULD\_NOT\_BE\_GRANTED (*new\_process*)

EXECUTION\_CLASS\_HAS\_NO\_USABLE\_EXECUTION\_SITES (execution class of *static\_context*)

EXECUTION\_SITE\_IS\_INACCESSIBLE (*site*)

EXECUTION\_SITE\_IS\_NOT\_IN\_EXECUTION\_CLASS (*site*, *static\_context*)

EXECUTION\_SITE\_IS\_UNKNOWN (*site*)

If *static\_context* is a foreign execution image:

FOREIGN\_EXECUTION\_IMAGE\_HAS\_NO\_SITE (*static\_context*)

FOREIGN\_SYSTEM\_IS\_INVALID (*site*, *new\_process*, HAS\_EXECUTIVE\_SYSTEM)

INTERPRETER\_IS\_INTERPRETABLE (interpreter of *static\_context*)



INTERPRETER\_IS\_NOT\_AVAILABLE (*static\_context*)  
 LABEL\_IS\_OUTSIDE\_RANGE (*new\_process*, the volume on which *new\_process* would reside)  
 LABEL\_IS\_OUTSIDE\_RANGE (*new\_process*, the would-be execution site of *new\_process*)  
 LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_PROCESSES\_PER\_USER)  
 REFERENCE\_CANNOT\_BE\_ALLOCATED  
 OBJECT\_OWNER\_VALUE\_WOULD\_BE\_INCONSISTENT\_WITH\_ATOMIC\_ACL  
 STATIC\_CONTEXT\_CONTENTS\_CANNOT\_BE\_EXECUTED (*static\_context*, *site*)  
 STATIC\_CONTEXT\_IS\_BEING\_WRITTEN (*static\_context*)  
 STATIC\_CONTEXT\_REQUIRES\_TOO\_MUCH\_MEMORY (*static\_context*)  
 If *process* is the calling process:  
     VOLUME\_IS\_FULL (calling process)

### 13.2.3 PROCESS\_GET\_WORKING\_SCHEMA

```

PROCESS_GET_WORKING_SCHEMA(
    process          : [ Process_designator ]
)
    sds_sequence    : Name_sequence
  
```

If no value is supplied for *process*, *process* designates the calling process.

PROCESS\_GET\_WORKING\_SCHEMA returns in *sds\_sequence* the sequence of SDS names of the SDSs forming the working schema of the process *process*.

If *process* is not the calling process a read lock of the default mode is established on *process*.

#### Errors

If *process* is not the calling process:

ACCESS\_ERRORS (*process*, ATOMIC, READ, READ\_LINKS)  
 PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, (READY, RUNNING, SUSPENDED, STOPPED))  
 PROCESS\_IS\_UNKNOWN (*process*)

### 13.2.4 PROCESS\_INTERRUPT\_OPERATION

```

PROCESS_INTERRUPT_OPERATION (
    process          : Process_designator
)
  
```

PROCESS\_INTERRUPT\_OPERATION interrupts a process.

There is no effect if *process* is not executing a PCTE operation; otherwise the interruption of all operations currently being executed by *process* is requested, with the following effect on *process*.

After a period of time, each such interrupted operation of *process*, whether suspended or not, is terminated with the error OPERATION\_IS\_INTERRUPTED. For any waiting operation the corresponding "process\_waiting\_for" link is deleted.

The time between the start of this operation and the end of the interruption of the operations of *process* is implementation-dependent.

#### Errors

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, WRITE\_ATTRIBUTES)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_EXECUTION)

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, RUNNING)

The following implementation-dependent error may be raised:

PROCESS\_IS\_THE\_CALLER (*process*)

PROCESS\_IS\_UNKNOWN (*process*)

NOTE - This operation is intended to provide the means for a tool to control other tools, e.g. to cause an operation of a deadlocked tool to be abandoned, or to interrupt a tool which is not itself controlling the duration of operations.

### 13.2.5 PROCESS\_RESUME

```
PROCESS_RESUME (
    process : Process_designator
)
```

PROCESS\_RESUME resumes the suspended *process process*, by changing its status to RUNNING.

#### Errors

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, WRITE\_ATTRIBUTES)

If the execution site of *process* is a foreign system:

FOREIGN\_SYSTEM\_IS\_INVALID (execution site of *process*, *process*,  
SUPPORTS\_IPC\_AND\_CONTROL)

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, SUSPENDED)

PROCESS\_IS\_THE\_CALLER (*process*)

PROCESS\_IS\_UNKNOWN (*process*)

### 13.2.6 PROCESS\_SET\_ALARM

```
PROCESS_SET_ALARM (
    duration : Natural
)
```

PROCESS\_SET\_ALARM changes the time left until alarm of the calling process to *duration*.

#### Errors

None.

### 13.2.7 PROCESS\_SET\_FILE\_SIZE\_LIMIT

```
PROCESS_SET_FILE_SIZE_LIMIT (
    process    : [ Process_designator ],
    fslimit    : Natural
)
```

If no value is supplied for *process*, *process* designates the calling process.

PROCESS\_SET\_FILE\_SIZE\_LIMIT changes the process file size limit of *process* to *fslimit*.

#### Errors

If *process* is not the calling process:

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, WRITE\_ATTRIBUTES)

If *fslimit* is greater than the current value of the process file size limit of *process*:

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_EXECUTION)

PROCESS\_IS\_UNKNOWN (*process*)

### 13.2.8 PROCESS\_SET\_OPERATION\_TIME\_OUT

```
PROCESS_SET_OPERATION_TIME_OUT (
    duration    : Natural
)
```

PROCESS\_SET\_OPERATION\_TIME\_OUT sets the process time-out of the calling process to *duration*.

#### Errors

None.

### 13.2.9 PROCESS\_SET\_PRIORITY

```
PROCESS_SET_PRIORITY (
    process    : [ Process_designator ],
    priority    : Natural
)
```

If no value is supplied for *process*, *process* designates the calling process.

PROCESS\_SET\_PRIORITY sets the process priority of *process* to MAX\_PRIORITY\_VALUE if *priority* is greater than MAX\_PRIORITY\_VALUE, and to *priority* otherwise.

#### Errors

If *process* is not the calling process:

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, WRITE\_ATTRIBUTES)

If *priority* is greater than the current value of the process priority of *process*:

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_EXECUTION)

PROCESS\_IS\_UNKNOWN (*process*)

### 13.2.10 PROCESS\_SET\_REFERENCED\_OBJECT

```
PROCESS_SET_REFERENCED_OBJECT (
    process      : [ Process_designator ],
    reference_name : Actual_key,
    object       : Object_designator
)
```

If no value is supplied for *process*, *process* designates the calling process.

PROCESS\_SET\_REFERENCED\_OBJECT sets a referenced object of *process* to *object*.

If a "referenced\_object" link from *process* with the key *reference\_name* already exists, its destination is changed to *object*. Otherwise, a "referenced\_object" link from *process* to *object* with the key *reference\_name* is created.

#### Errors

If *process* is not the calling process:

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, APPEND\_LINKS)

If *process* is not the calling process and there is a referenced object *reference\_name*:

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, WRITE\_LINKS)

If *process* is not the calling process:

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, READY)

PROCESS\_IS\_UNKNOWN (*process*)

REFERENCE\_NAME\_IS\_INVALID (*reference\_name*)

REFERENCED\_OBJECT\_IS\_NOT\_MUTABLE (*reference\_name*)

If *process* is the calling process:

VOLUME\_IS\_FULL (calling process)

### 13.2.11 PROCESS\_SET\_TERMINATION\_STATUS

```
PROCESS_SET_TERMINATION_STATUS (
    termination_status : Integer
)
```

PROCESS\_SET\_TERMINATION\_STATUS provides a value *termination\_status* to be stored in the process termination status of the calling process when it terminates, provided it is not terminated by PROCESS\_TERMINATE with a termination status parameter.

#### Errors

VOLUME\_IS\_FULL (calling process)

### 13.2.12 PROCESS\_SET\_WORKING\_SCHEMA

```
PROCESS_SET_WORKING_SCHEMA (
    process      : [ Process_designator ],
    sds_sequence : Name_sequence
)
```

If no value is supplied for *process*, *process* designates the calling process.

PROCESS\_SET\_WORKING\_SCHEMA sets the working schema of a process according to the sequence of SDSs identified by the SDS names in *sds\_sequence*, replacing the current working schema, if there is one.

If *process* is the calling process:

- the previous "sds\_in\_working\_schema" links of *process* and the "in\_working\_schema\_of" links from each previous SDS in working schema to *process* are deleted;
- for each SDS *sds* identified by *sds\_sequence*(*I*), an "sds\_in\_working\_schema" link with key *I* (starting from *I* = 1) from *process* to *sds* and an "in\_working\_schema\_of" link from *sds* to *process* are created.

If *process* is not the calling process (and *process* is ready):

- the previous "sds\_in\_working\_schema" links of *process* are deleted;
- for each SDS *sds* identified *sds\_sequence*(*I*), an "sds\_in\_working\_schema" link with key *I* (starting from *I* = 1) from *process* to *sds* is created.

A new working schema is created as follows:

- The sequence of SDS names is set to *sds\_sequence*.
- The set of types in working schema is constituted as follows:
  - a type in working schema is created for each type associated with a type in SDS in an SDS of *sds\_sequence*;
  - the types in SDS of each created type in working schema are set to the types in SDS with the same associated type, and the composite names of those types in SDS;
  - the usage mode of each created type in working schema is set to the union of the usage modes of all its types in SDS;
  - the other properties of the created types in working schema are determined from their types in SDS (see 8.5).

## Errors

If *process* is not the calling process and an "sds\_in\_working\_schema" link exists:

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, WRITE\_LINKS)

If *process* is not the calling process:

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (SDS with name in *sds\_sequence*, ATOMIC, SYSTEM\_ACCESS)

If *process* is the calling process:

DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED (SDS with name in *sds\_sequence*, ATOMIC, EXPLOIT\_SCHEMA)

If *process* is not the calling process:

DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED\_TO\_PROCESS (*process*, SDS with name in *sds\_sequence*, ATOMIC, EXPLOIT\_SCHEMA)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_SDS\_IN\_WORKING\_SCHEMA)

If *process* is not the calling process:

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, READY)

PROCESS\_IS\_UNKNOWN (*process*)

If *process* is the calling process:

SDS\_IS\_UNDER\_MODIFICATION (SDS with name in *sds\_sequence*)

SDS\_IS\_UNKNOWN (SDS with name in *sds\_sequence*)

SDS\_WOULD\_APPEAR\_TWICE\_IN\_WORKING\_SCHEMA (*sds\_sequence*)

If *process* is the calling process:

VOLUME\_IS\_FULL (calling process)

#### NOTES

1 A process need not have the predefined SDSs in its working schema in order to call operations except operations defined in clause 9 operating on objects or links with types and types in SDS defined in the predefined SDSs.

2 Setting the working schema is independent of activities. In order to maintain the integrity of working schemas, operations which affect the typing information contained in SDSs included in a working schema are explicitly prohibited and a working schema which contains an SDS with uncommitted modifications of the typing information may not be created.

### 13.2.13 PROCESS\_START

```
PROCESS_START (
    process      : Process_designator,
    arguments    : String,
    environment  : String,
    site         : [ Execution_site_designator ],
    initial_status : Initial_status
)
```

PROCESS\_START starts the execution of the static context or foreign execution image *static\_context* of a process that has already been created.

If *site* is supplied the "executed\_on" link of *process* is replaced (if different) by one with destination *site*; otherwise *site* is the destination of that link

The process status of *process* is changed to *initial\_status*, provided *process* is ready.

A link is created to *process* from each destination of the following links:

- "in\_working\_schema\_of" from each destination of "sds\_in\_working\_schema";
- "running\_process" from the destination of "executed\_on";
- "opened\_by" from each destination of "open\_object";
- "process\_started\_in" from the destination of "started\_in\_activity";
- "user\_identity\_of" from the destination of "user\_identity";
- "adopted\_user\_group\_of" from the destination of "adopted\_user\_group";
- "consumer\_process" from the destination of "consumer\_identity", if any.

These links are created even if any origin object is on a read-only volume or is a replicated copy.

For each "open\_object" link of *process*, the contents of the destination are opened and the current position in the object contents is shared with the process that created *process*. Furthermore, if the parent process of *process* is not the calling process, then:

- a lock is acquired by the activity of the parent process on the opened object;
- if the calling process has closed the object contents, then the current position of the opened object is determined in the same way as by CONTENTS\_OPEN, in the opening mode of the "open\_object" link.

If *static\_context* is interpretable, an "actual\_interpreter" link is created to the interpreter of the interpretable static context, if it has one, else to the default interpreter of *process* if it has one, else to the default interpreter of the home object of *process*, provided there is a home object and it has a default interpreter.

The "process\_string\_arguments" and "process\_environment" attributes of *process* are set to *arguments* and *environment* respectively.

### Errors

ACCESS\_ERRORS (*static\_context*, ATOMIC, READ, EXECUTE)

ACCESS\_ERRORS (any interpreter, ATOMIC, READ, EXECUTE)

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, WRITE\_ATTRIBUTES)

For each SDS *sds* which is the destination of an "in\_working\_schema\_of" link from *process*:

ACCESS\_ERRORS (*sds*, ATOMIC, SYSTEM\_ACCESS)

DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED\_TO\_PROCESS (*process*,  
*sds*, ATOMIC, EXPLOIT\_SCHEMA)

SDS\_IS\_UNDER\_MODIFICATION (*sds*)

SDS\_IS\_UNKNOWN (*sds*)

For the execution site *site* which is the destination of an "executed\_on" link from *process*:

ACCESS\_ERRORS (*site*, ATOMIC, SYSTEM\_ACCESS)

For each open object *object* which is the destination of an "open\_object" link from *process*:

ACCESS\_ERRORS (*object*, ATOMIC, SYSTEM\_ACCESS)

If the link's opening mode attribute is READ\_ONLY or READ\_WRITE,  
DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED\_TO\_PROCESS (*process*,  
*object*, ATOMIC, READ, READ\_CONTENTS)

If the link's opening mode attribute is WRITE\_ONLY or READ\_WRITE,  
DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED\_TO\_PROCESS (*process*,  
*object*, ATOMIC, READ, WRITE\_CONTENTS)

If the link's opening mode attribute is APPEND\_ONLY,  
DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED\_TO\_PROCESS (*process*,  
*object*, ATOMIC, READ, APPEND\_CONTENTS)

For the activity *activity* which is the destination of the "started\_in\_activity" link from *process*:

ACCESS\_ERRORS (*activity*, ATOMIC, SYSTEM\_ACCESS)

ACTIVITY\_STATUS\_IS\_INVALID (*activity*, ACTIVE)

For the user *user* which is the destination of the "user\_identity" link from *process*:

ACCESS\_ERRORS (*user*, ATOMIC, SYSTEM\_ACCESS)

For the user group *group* which is the destination of the "adopted\_user\_group" link from *process*:

ACCESS\_ERRORS (*group*, ATOMIC, SYSTEM\_ACCESS)



For the consumer group *group* which is the destination of the "consumer\_identity" link from *process*:

ACCESS\_ERRORS (*group*, ATOMIC, SYSTEM\_ACCESS)  
 ACCESS\_ERRORS (parent process of *process*, ATOMIC, SYSTEM\_ACCESS)  
 EXECUTION\_SITE\_IS\_INACCESSIBLE (*site*)  
 EXECUTION\_SITE\_IS\_NOT\_IN\_EXECUTION\_CLASS (*site*, *static\_context*)  
 EXECUTION\_SITE\_IS\_UNKNOWN (*site*)

If *site* is a foreign system:

FOREIGN\_SYSTEM\_IS\_INVALID (*site*, *process*, (HAS\_EXECUTIVE\_SYSTEM,  
 SUPPORTS\_EXECUTIVE\_CONTROL, SUPPORTS\_MONITOR))  
 INTERPRETER\_IS\_INTERPRETABLE (interpreter of *static\_context*)  
 INTERPRETER\_IS\_NOT\_AVAILABLE (*static\_context*)  
 LABEL\_IS\_OUTSIDE\_RANGE (*process*, *site*)  
 LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_PROCESSES)  
 PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, READY)  
 PROCESS\_IS\_UNKNOWN (*process*)  
 STATIC\_CONTEXT\_CONTENTS\_CANNOT\_BE\_EXECUTED (*static\_context*, *site*)  
 STATIC\_CONTEXT\_IS\_BEING\_WRITTEN (*static\_context*)  
 STATIC\_CONTEXT\_REQUIRES\_TOO\_MUCH\_MEMORY (*static\_context*)

### 13.2.14 PROCESS\_SUSPEND

```
PROCESS_SUSPEND (
    process : [ Process_designator ],
    alarm   : [ Natural ]
)
```

If no value is supplied for *process*, *process* designates the calling process.

PROCESS\_SUSPEND suspends a running process.

PROCESS\_SUSPEND changes the status of *process* to SUSPENDED, provided it already has the value RUNNING; and if *alarm* is supplied and *process* is the calling process sets the value of time left until alarm to *alarm*.

If time left until alarm is non-zero, it defines a maximum duration in seconds for the suspension of the process. If this duration expires, the process receives an implementation-dependent alarm message of message type WAKE (provided it has reserved a message queue and is the listened-to process for the message queue) and is resumed.

#### Errors

If *process* is not the calling process:

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, WRITE\_ATTRIBUTES)

If the execution site of *process* is a foreign system:

FOREIGN\_SYSTEM\_IS\_INVALID (execution site of *process*, *process*,  
 SUPPORTS\_IPC\_AND\_CONTROL)

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, RUNNING)

If *alarm* is supplied:

PROCESS\_IS\_NOT\_THE\_CALLER (*process*)

PROCESS\_IS\_UNKNOWN (*process*)

### 13.2.15 PROCESS\_TERMINATE

```
PROCESS_TERMINATE (
    process          : [ Process_designator ],
    termination_status : [ Integer ]
)
```

If no value is supplied for *process*, *process* designates the calling process.

PROCESS\_TERMINATE terminates a process. In certain conditions this results in the (composite) deletion of the process or its parent.

Any ongoing operations invoked from *process* are interrupted in the same way as by PROCESS\_INTERRUPT\_OPERATION.

PROCESS\_TERMINATE changes the links and attributes of *process* as follows:

- "process\_status" is set to TERMINATED;
- "process\_termination\_time" is set to the current time;
- "process\_termination\_status" is set to *termination\_status*, if supplied, otherwise to EXIT\_ERROR if *process* is the calling process or FORCED\_TERMINATION if not.

Destinations of "open\_object" links from *process* are closed.

The following links from *process* and their reverse links are deleted:

- "sds\_in\_working\_schema" links and their reverse "in\_working\_schema\_of" links;
- "executed\_on" links and their reverse "running\_process" links;
- "opened\_object" links and their reverse "opened\_by" links;
- "reserved\_message\_queue" links and their reverse "reserved\_by" links;
- "user\_identity" links and their reverse "user\_identity\_of" links;
- "adopted\_user\_group" links and their reverse "adopted\_user\_group\_of" links;
- "consumer\_identity" links and their reverse "consumer\_process" links.

The "adoptable\_user\_group" links from *process* are deleted.

If the parent of *process* is waiting for termination of *process* then the parent of *process* discontinues waiting and "acknowledged\_termination" of *process* is set to **true**.

If "deletion\_upon\_termination" and "acknowledged\_termination" of *process* are both **true**, all component processes of *process* are ready or terminated, and the conditions for the object deletion of *process* hold (see 9.3.5), then *process* is deleted.

If the parent of *process* is terminated, "deletion\_upon\_termination" and "acknowledged\_termination" of the parent of *process* are both **true**, all component processes of the parent of *process* are ready or terminated, and the conditions for the object deletion of the parent of *process* hold (see 9.3.5), then the parent of *process* is deleted.

If deletion of a "process" object is not possible, then the "child\_process" link remains.

If an activity is initiated by a process and no corresponding ACTIVITY\_ABORT or ACTIVITY\_END call is made before termination of the process, then an ACTIVITY\_END call is implied if the termination status of the terminating process is EXIT\_SUCCESS and an ACTIVITY\_ABORT call is implied otherwise.

When a process is started, and "delete\_upon\_termination" is **true**, then the activity in which the process starts acquires a default delete lock on the process. If "delete\_upon\_termination" is **false** when the process starts and is set **true** while the process is running then the delete lock is acquired at that point.

### Errors

If *process* is not the calling process:

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, (WRITE\_ATTRIBUTES, WRITE\_LINKS))

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, (READY, RUNNING, STOPPED, SUSPENDED))

PROCESS\_IS\_INITIAL\_PROCESS (*process*)

PROCESS\_IS\_UNKNOWN (*process*)

### 13.2.16 PROCESS\_UNSET\_REFERENCED\_OBJECT

```
PROCESS_UNSET_REFERENCED_OBJECT (
    process      : [ Process_designator ],
    reference_name : Actual_key
)
```

If no value is supplied for *process*, *process* designates the calling process.

PROCESS\_UNSET\_REFERENCED\_OBJECT unsets a referenced object of *process*.

If there is no "referenced\_object" link from *process* with key *reference\_name*, the operation has no effect. Otherwise, the "referenced\_object" link from *process* with key *reference\_name* is deleted.

### Errors

If *process* is not the calling process and there is a "referenced\_object" link from *process* with key *reference\_name*:

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, WRITE\_LINKS)

If *process* is not the calling process:

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, READY)

PROCESS\_IS\_UNKNOWN (*process*)

REFERENCE\_NAME\_IS\_INVALID (*reference\_name*)

REFERENCED\_OBJECT\_IS\_NOT\_MUTABLE (*reference\_name*)

### 13.2.17 PROCESS\_WAIT\_FOR\_ANY\_CHILD

```
PROCESS_WAIT_FOR_ANY_CHILD (
)
    termination_status : Integer,
    child                : Natural
```

PROCESS\_WAIT\_FOR\_ANY\_CHILD sets the calling thread of the calling process waiting until any of its child processes has terminated.

If any child of the calling process has process status TERMINATED and acknowledged termination **false**, the acknowledged termination of the terminated child process is set to **true** and that child process is deleted if it has deletion upon termination **true** and the deletion conditions (see 13.2.15) are satisfied. If more than one child process fulfils the condition, one is selected in an implementation-defined manner to fill the role of terminated child process.

If no child of the calling process has process status TERMINATED, the operation waits and a "process\_waiting\_for" link is created to the calling process with waiting type WAITING\_FOR\_TERMINATION. The operation continues when the process status of any child process changes to TERMINATED.

The process termination status of the terminated child is returned in *termination\_status*, unless the confidentiality label of the calling process does not dominate that of the terminated child process or the integrity label of the calling process is not dominated by that of the terminated child process, in which cases *termination\_status* is set to UNAVAILABLE (a binding-defined value different from the named values of the "process\_termination\_status" attribute). The key of the "child\_process" link from the calling process to the terminated child is returned in *child*.

#### Errors

DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED (the terminated child process, ATOMIC, WRITE\_ATTRIBUTES)

PROCESS\_HAS\_NO\_UNTERMINATED\_CHILD

### 13.2.18 PROCESS\_WAIT\_FOR\_CHILD

```
PROCESS_WAIT_FOR_CHILD (
    child                : Process_designator
)
    termination_status : Integer
```

PROCESS\_WAIT\_FOR\_CHILD sets the calling thread of the calling process waiting until the nominated child process has terminated.

If *child* has process status TERMINATED and acknowledged termination **false**, the acknowledged termination of *child* is set to **true** and *child* is deleted if it has deletion upon termination **true** and the deletion conditions (see 13.2.15) are satisfied.

Otherwise, the operation waits and a "process\_waiting\_for" link is created to *child* with waiting type WAITING\_FOR\_TERMINATION. The operation continues when the process status of *child* changes to TERMINATED.

PROCESS\_WAIT\_FOR\_CHILD returns the process termination status of *child* in *termination\_status*, unless the confidentiality label of the calling process does not dominate that

of *child* or the integrity label of the calling process is not dominated by that of *child*, in which cases *termination\_status* is set to UNAVAILABLE (a binding-defined value different from the named values of the "process\_termination\_status" attribute).

### Errors

DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED (*child*, ATOMIC, WRITE\_ATTRIBUTES)

PROCESS\_IS\_NOT\_TERMINABLE\_CHILD (*child*)

PROCESS\_IS\_UNKNOWN (*child*)

PROCESS\_TERMINATION\_IS\_ALREADY\_ACKNOWLEDGED (*child*)

## 13.3 Security operations

### 13.3.1 PROCESS\_ADOPT\_USER\_GROUP

```
PROCESS_ADOPT_USER_GROUP (
    process      : [ Process_designator ],
    user_group   : User_group_designator
)
```

If no value is supplied for *process*, *process* designates the calling process.

PROCESS\_ADOPT\_USER\_GROUP changes the adopted user group of *process* to *user\_group*.

Let *G* be *user\_group*, *P* be *process*, and *G'* be the previous adopted user group.

If *process* is the calling process:

- The following links are deleted:
  - . "adopted\_user\_group" from *P* to *G'*;
  - . "adopted\_user\_group\_of" from *G'* to *P*;
  - . "adoptable\_user\_group" from *P* to *G*.
- The following links are created, setting the key values to the next available natural in each case:
  - . "adopted\_user\_group" from *P* to *G*;
  - . "adopted\_user\_group\_of" from *G* to *P*;
  - . "adoptable\_user\_group" from *P* to *G*.
- The effective security groups of the process are changed to consist of:
  - . The user *user* of *process* (no change);
  - . *user\_group*;
  - . all the supergroups of *user\_group*;
  - . all the program groups to which the static contexts run or executed by the calling process belong (no change);
  - . all the supergroups to which these program groups belong (no change).

If *process* is not the calling process (and is ready):

- The following links are deleted:
  - . "adopted\_user\_group" from P to G';
  - . "adoptable\_user\_group" from P to G.
- The following links are created, setting the key values to the next available natural in each case:
  - . "adopted\_user\_group" from P to G;
  - . "adoptable\_user\_group" from P to G'.

If P is the calling process, there is a "consumer\_identity" link from P to a consumer group object C, and if the new effective security groups are such that EVALUATE\_PROCESS (P, C, EXPLOIT\_CONSUMER\_IDENTITY) is **false** (see 19.1.2), then the "consumer\_identity" link from P to C and the "consumer\_process" link from C to P are deleted.

The working schema of *process* is reset to empty by deleting all "sds\_in\_working\_schema" links from *process* and their reverse "in\_working\_schema\_of" links.

### Errors

Access errors are determined on the basis of the discretionary context in force before the change in the effective security groups which this operation produces.

If *process* is not the calling process:

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, APPEND\_LINKS)

If *process* is not the calling process:

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, WRITE\_LINKS)

OBJECT\_IS\_INACCESSIBLE (*user\_group*, ATOMIC)

OBJECT\_IS\_INACCESSIBLE (G', ATOMIC)

If *process* is not the calling process:

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, READY)

PROCESS\_IS\_UNKNOWN (*process*)

SECURITY\_GROUP\_IS\_NOT\_ADOPTABLE (*user\_group*)

SECURITY\_GROUP\_IS\_UNKNOWN (*user\_group*)

USER\_IS\_NOT\_MEMBER (*user*, *user\_group*)

If *process* is the calling process:

VOLUME\_IS\_FULL (calling process)

### NOTES

1 This operation changes the user group which is currently adopted by the designated process, and therefore changes the role in which the user is acting.

2 Users may be removed from user groups at any time. It is therefore necessary to check that the current user is still a member of the designated user group before adopting it. It is insufficient to rely on the "adoptable\_user\_group" links.

**13.3.2 PROCESS\_GET\_DEFAULT\_ACL**

```
PROCESS_GET_DEFAULT_ACL (
)
    acl : Acl
```

PROCESS\_GET\_DEFAULT\_ACL returns the default atomic ACL of the calling process as *acl*.

**Errors**

None.

**13.3.3 PROCESS\_GET\_DEFAULT\_OWNER**

```
PROCESS_GET_DEFAULT_OWNER (
)
    group : Group_identifier
```

PROCESS\_GET\_DEFAULT\_OWNER returns the group identifier of the default object owner of the calling process as *group*.

**Errors**

None.

**13.3.4 PROCESS\_SET\_ADOPTABLE\_FOR\_CHILD**

```
PROCESS_SET_ADOPTABLE_FOR_CHILD (
    process      : [ Process_designator ],
    user_group   : User_group_designator,
    adoptability : Boolean
)
```

If no value is supplied for *process*, *process* designates the calling process.

PROCESS\_SET\_ADOPTABLE\_FOR\_CHILD changes the "adoptable\_for\_child" attribute of the "adoptable\_user\_group" link from *process* to *user\_group* to *adoptability*.

**Errors**

If *process* is not the calling process:

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, WRITE\_LINKS)

If *process* is not the calling process:

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, READY)

PROCESS\_IS\_UNKNOWN (*process*)

SECURITY\_GROUP\_IS\_UNKNOWN (*user\_group*)

SECURITY\_GROUP\_IS\_NOT\_ADOPTABLE (*user\_group*, *process*)

**13.3.5 PROCESS\_SET\_DEFAULT\_ACL\_ENTRY**

```
PROCESS_SET_DEFAULT_ACL_ENTRY (
    process      : [ Process_designator ],
    group        : Group_identifier,
    modes        : Atomic_access_rights
)
```



If no value is supplied for *process*, *process* designates the calling process.

PROCESS\_SET\_DEFAULT\_ACL\_ENTRY changes the default atomic ACL of the process *process*.

The ACL entry for *group* in the "default\_atomic\_acl" attribute of *process* is set to *modes*.

#### Errors

If *process* is not the calling process:

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, WRITE\_ATTRIBUTES)

DEFAULT\_ACL\_WOULD\_BE\_INVALID (*process*, *group*, *modes*)

DEFAULT\_ACL\_WOULD\_BE\_INCONSISTENT\_WITH\_DEFAULT\_OBJECT\_OWNER (*process*, *group*)

If *process* is not the calling process:

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, READY)

PROCESS\_IS\_UNKNOWN (*process*)

SECURITY\_GROUP\_IS\_UNKNOWN (*group*)

### 13.3.6 PROCESS\_SET\_DEFAULT\_OWNER

```
PROCESS_SET_DEFAULT_OWNER (
    process : [ Process_designator ],
    group   : Group_identifier
)
```

If no value is supplied for *process*, *process* designates the calling process.

PROCESS\_SET\_DEFAULT\_OWNER changes the default object owner of *process* to the security group identifier *group*.

#### Errors

If *process* is not the calling process:

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, WRITE\_ATTRIBUTES)

DEFAULT\_ACL\_WOULD\_BE\_INCONSISTENT\_WITH\_DEFAULT\_OBJECT\_OWNER (*process*, *group*)

PROCESS\_IS\_UNKNOWN (*process*)

If *process* is not the calling process:

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, READY)

SECURITY\_GROUP\_IS\_UNKNOWN (*group*)

### 13.3.7 PROCESS\_SET\_USER

```
PROCESS_SET_USER (
    user       : User_designator,
    user_group : User_group_designator
)
```

PROCESS\_SET\_USER sets the user of the calling process to *user* and changes the adopted user group of the calling process to *user\_group*.

Let P be the calling process, U be the previous user of the process, G be the previous adopted user group, U' be *user*, and G' be *user\_group*.

The following links are deleted:

- "user\_identity" from P to U;
- "user\_identity\_of" from U to P;
- "adopted\_user\_group" from P to G;
- "adopted\_user\_group\_of" from G to P;
- "adoptable\_user\_group" from P to the set of user groups currently so linked, excluding G.

The following links are created, setting the key values to the next available integer in each case:

- "user\_identity" from P to U';
- "user\_identity\_of" from U' to P;
- "adopted\_user\_group" from P to G';
- "adopted\_user\_group\_of" from G' to P;
- "adoptable\_user\_group" from P to each user group of which U' is a member, excluding G'.

The effective security groups of the process are changed to consist of:

- *user*;
- *user\_group*;
- all the supergroups of *user\_group*;
- all the program groups to which the static contexts run or executed by the calling process belong, unchanged;
- all the supergroups to which these program groups belong (no change).

Let W be the execution site of P and V be the volume on which P resides, and let P' be P as updated by the operation. The confidentiality label of P is set to the confidentiality label C which is the conjunction of confidentiality low label of W and the confidentiality low label of V, providing that the following are all **true** (see 20.1.3, 20.1.4):

LABEL\_DOMINATES (confidentiality clearance of *user*, C)

CONFIDENTIALITY\_LABEL\_WITHIN\_RANGE (P', W)

CONFIDENTIALITY\_LABEL\_WITHIN\_RANGE (P', V)

The integrity context of P is set to the integrity label I which is the disjunction of the user's integrity clearance, the integrity high label of W, and the integrity high label of V, providing that the following are all **true**:

LABEL\_DOMINATES (I, integrity clearance of *user*)

INTEGRITY\_LABEL\_WITHIN\_RANGE (P', W)

INTEGRITY\_LABEL\_WITHIN\_RANGE (P', V)

If there is a link of type "consumer\_identity" from P to a consumer group C, and if the new effective security groups are such that

EVALUATE\_PROCESS (P, C, EXPLOIT\_CONSUMER\_IDENTITY)

is **false**, then the "consumer\_identity" link from P to C and the "consumer\_process" link from C to P are deleted.

The working schema of the calling process is reset to its initial value (empty) by deleting all "sds\_in\_working\_schema" links from the calling process and their reverse "in\_working\_schema\_of" links.

### Errors

Error conditions are determined on the basis of the discretionary context in force before the change in the effective security groups which this operation produces.

ACCESS\_ERRORS (*user*, ATOMIC, SYSTEM\_ACCESS)

ACCESS\_ERRORS (*user\_group*, ATOMIC, SYSTEM\_ACCESS)

OBJECT\_IS\_INACCESSIBLE (U, ATOMIC)

OBJECT\_IS\_INACCESSIBLE (G, ATOMIC)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SECURITY)

PROCESS\_LABELS\_WOULD\_BE\_INCOMPATIBLE (*user*)

SECURITY\_GROUP\_IS\_UNKNOWN (*user*)

SECURITY\_GROUP\_IS\_UNKNOWN (*user\_group*)

USER\_IS\_NOT\_MEMBER (*user*, *user\_group*)

NOTE - This operation establishes the user on behalf of which the current process will run, and the role in which the user will act. It is intended to be used by the user authentication tool.

## 13.4 Profiling operations

### 13.4.1 PROCESS\_PROFILING\_OFF

```
PROCESS_PROFILING_OFF (
    handle : Profile_handle
)
    buffer : Buffer
```

PROCESS\_PROFILING\_OFF terminates the profiling of the calling process initiated with the profile handle *handle*, and returns the results in *buffer*.

### Errors

PROFILING\_IS\_NOT\_SWITCHED\_ON (*handle*)

### 13.4.2 PROCESS\_PROFILING\_ON

```
PROCESS_PROFILING_ON (
    start   : Address,
    end     : Address,
    count   : Natural
)
    handle : Profile_handle
```

PROCESS\_PROFILING\_ON initiates profiling of the calling process. Profiling is an implementation-defined action. It continues until the operation PROCESS\_PROFILING\_OFF is called with the returned profile handle *handle* or the calling process terminates.

If profiling is already initiated for the calling process, it is reinitiated with a new profiling buffer identified by the returned profile handle.

### Errors

MEMORY\_REGION\_IS\_NOT\_IN\_PROFILING\_SPACE (*start*, *end*)

NOTE - Profiling is implementation-defined but is intended to provide in each element of a profiling buffer identified by *handle* a count of the number of times the process was executing at, or accessing, a memory address associated with that element. *start* and *end* specify a region of the process memory to be profiled: the mapping to elements of the buffer is implementation-defined. Other calls of PROCESS\_PROFILING\_ON can request profiling into other buffers.

## 13.5 Monitoring operations

### 13.5.1 PROCESS\_ADD\_BREAKPOINT

```
PROCESS_ADD_BREAKPOINT (
    process : Process_designator,
    breakpoint : Address
)
```

PROCESS\_ADD\_BREAKPOINT adds a breakpoint for *process*. The effect is implementation-defined.

### Errors

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, WRITE\_CONTENTS)

MEMORY\_ADDRESS\_IS\_OUT\_OF\_PROCESS (*breakpoint*, *process*)

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, STOPPED)

PROCESS\_IS\_NOT\_CHILD (*process*)

PROCESS\_IS\_UNKNOWN (*process*)

NOTE - The format of a breakpoint is implementation-defined but it is intended to define an instruction or data address, access to which will cause the accessing thread of *process* to stop.

### 13.5.2 PROCESS\_CONTINUE

```
PROCESS_CONTINUE (
    process : Process_designator
)
```

PROCESS\_CONTINUE continues any stopped threads of a process.

The status of *process* is set to RUNNING.

### Errors

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, WRITE\_CONTENTS)

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, STOPPED)

PROCESS\_IS\_NOT\_CHILD (*process*)

PROCESS\_IS\_UNKNOWN (*process*)

### 13.5.3 PROCESS\_PEEK

```
PROCESS_PEEK (
    process    : Process_designator,
    address    : Address
)
    value      : Process_data
```

PROCESS\_PEEK returns as *value* the contents at *address* of *process*.

#### Errors

ACCESS\_ERRORS (*process*, ATOMIC, READ, READ\_CONTENTS)

MEMORY\_ADDRESS\_IS\_OUT\_OF\_PROCESS (*address*, *process*)

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, STOPPED)

PROCESS\_IS\_NOT\_CHILD (*process*)

PROCESS\_IS\_UNKNOWN (*process*)

### 13.5.4 PROCESS\_POKE

```
PROCESS_POKE (
    process    : Process_designator,
    address    : Address,
    value      : Process_data
)
```

PROCESS\_POKE modifies *process* at *address* to *value*.

#### Errors

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, WRITE\_CONTENTS)

MEMORY\_ADDRESS\_IS\_OUT\_OF\_PROCESS (*address*, *process*)

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, STOPPED)

PROCESS\_IS\_NOT\_CHILD (*process*)

PROCESS\_IS\_UNKNOWN (*process*)

### 13.5.5 PROCESS\_REMOVE\_BREAKPOINT

```
PROCESS_REMOVE_BREAKPOINT (
    process    : Process_designator,
    breakpoint : Address
)
```

PROCESS\_REMOVE\_BREAKPOINT removes a breakpoint *breakpoint* of process *process*.

#### Errors

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, WRITE\_CONTENTS)

BREAKPOINT\_IS\_NOT\_DEFINED (*breakpoint*)

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, STOPPED)

PROCESS\_IS\_NOT\_CHILD (*process*)

PROCESS\_IS\_UNKNOWN (*process*)

### 13.5.6 PROCESS\_WAIT\_FOR\_BREAKPOINT

```
PROCESS_WAIT_FOR_BREAKPOINT (
    process : Process_designator
)
    breakpoint : Address
```

PROCESS\_WAIT\_FOR\_BREAKPOINT sets the calling thread of the calling process waiting until the process *process* is stopped or terminated.

If *process* has process status TERMINATED, the error condition PROCESS\_LACKS\_REQUIRED\_STATUS occurs.

Otherwise, a "process\_waiting\_for" link is created to *process* with waiting type WAITING\_FOR\_TERMINATION. The operation waits until *process* reaches a breakpoint, in which case the breakpoint is returned in *breakpoint*, or the process status of *process* changes to TERMINATED, in which case the error condition PROCESS\_IS\_TERMINATED occurs.

#### Errors

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, READ\_CONTENTS)

PROCESS\_IS\_NOT\_CHILD (*process*)

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, (READY, RUNNING, STOPPED, SUSPENDED))

PROCESS\_IS\_UNKNOWN (*process*)

## 14 Message queues

### 14.1 Message queue concepts

```
Message ::
    DATA : seq of Octet
    MESSAGE_TYPE : Message_type
```

```
Received_message ::
    MESSAGE : Message
    POSITION : Natural
```

```
Message_type = Standard_message_type | Notification_message_type
              | Implementation_defined_message_type | Undefined_message_type
```

```
Message_types = set of Message_type | ALL_MESSAGE_TYPES
```

```
Standard_message_type = interrupt | quit | finish | suspend | end | abort |
                        deadlock | wake
```

```
Implementation_defined_message_type :: Token
```

```
Undefined_message_type :: Token
```

```
Handler :: Token
```

**sds** system:

message\_queue: **child type of object with attribute**

reader\_waiting: (read) **non\_duplicated boolean**;  
 writer\_waiting: (read) **non\_duplicated boolean**;  
 space\_used: (read) **non\_duplicated natural**;  
 total\_space: (read) **natural**;  
 message\_count: (read) **non\_duplicated natural**;  
 last\_send\_time: (read) **non\_duplicated time**;  
 last\_receive\_time: (read) **non\_duplicated time**;

**link**

reserved\_by: (navigate) **non\_duplicated designation link to process**;  
 listened\_to: (navigate) **non\_duplicated designation link to process**;  
 notifier: (navigate) **non\_duplicated designation link** (notifier\_key: **natural**) to  
 object with

**attribute**

modification\_event: (read) **boolean**;  
 change\_event: (read) **boolean**;  
 delete\_event: (read) **boolean**;  
 move\_event: (read) **boolean**;

**end** notifier;

**end** message\_queue;

**end** system;

Messages and message queues allow processes to communicate. A message queue has an associated sequence of messages. A message contains data, and has a message type. The space occupied by a message is implementation-defined. For notification message types see 15.1.2.

Implementation\_defined\_message\_type and Undefined\_message\_type are implementation-defined types disjoint from each other and from Standard\_message\_type and Notification\_message\_type. The meanings of implementation-defined message types are implementation-defined. For the intended meanings of standard message types see Note 3 below.

The value ALL\_MESSAGE\_TYPES denotes the set of all message types, including implementation-defined and undefined message types.

Each message in a message queue is assigned a *position number* which it retains while it is in the queue. The position numbers are positive naturals, monotonically increasing with time of arrival in the queue but otherwise implementation-dependent.

Reader waiting is **true** if and only if one or more processes are waiting to receive a message.

Writer waiting is **true** if and only if one or more processes are waiting to send a message.

The space used is the space currently required by the message queue to hold its messages, in octets.

The total space is the maximum possible size of the space used. This may vary between message queues and is initialized to an implementation-defined value which must not be less than four times MAX\_MESSAGE\_SIZE (see clause 24). An implementation may place an upper limit on the total space of a message queue; this must not be less than MAX\_MESSAGE\_QUEUE\_SPACE (see clause 24).

The message count is the number of messages in the message queue.



The last send time and last receive time record the system time on the last occasion that a message was sent to the queue and received from the queue, respectively. Initially they are equal to the default initial value for time attributes. If the last sent message was sent through MESSAGE\_SEND\_WAIT, the last send time is the system time when the message actually entered the queue (at the end of the waiting period). If the last received message was received through MESSAGE\_RECEIVE\_WAIT, last receive time is the system time when the message was actually received from the queue (at the end of the waiting period).

The destination of the "reserved\_by" link, if any, is called the *reserving* process of the message queue; it is also said to have *reserved* the message queue. The reserving process is the only process which can receive or peek messages from the message queue, and it must have adequate read access permission to the message queue. If there is no reserving process then any process can receive or peek messages, subject to access permission.

The destination of the "listened\_to" link, if any, is the reserving process (which must exist), and indicates that the reserving process has an associated procedure which is executed on the raising of a message queue event by the appearance of a message of one of a specified set of message types in the message queue. Such a procedure is called a *handler*. In this case the reserving process is called the *listening process* of the message queue. The "listened\_to" link is reversed by an "is\_listener" link, with an attribute which defines the message types of messages for which the handler is executed. The handler is invoked with a single argument, which denotes the affected message queue in a binding-defined manner. The types of a handler and of its argument are binding-defined.

Notifiers are described in 15.1.

#### NOTES

- 1 An implicit modification of predefined attributes of a message queue does not require a write lock on the message queue.
- 2 The intended meanings of the standard message types are as follows.
  - INTERRUPT: user interruption;
  - QUIT: user wants to quit;
  - FINISH: the receiving process should terminate;
  - SUSPEND: the receiving process should suspend itself;
  - END: the current activity of the receiving process should be normally terminated;
  - ABORT: the current activity of the receiving process should be abnormally terminated;
  - DEADLOCK: deadlock has been detected;
  - WAKE: the receiving process's time left until alarm has expired (see 13.1.4).
- 3 A process can send messages to itself. A message queue can have several concurrent readers if there is no reserving process. A message queue can have several concurrent writers. If more than one process is eligible to receive a message, it is not defined which of the eligible processes receives it.
- 4 The associated sequences of messages of message queues are not affected by transaction rollback.
- 5 The handler is executed and the listening process status is changed to RUNNING, even if the listening process is suspended — see 14.1.

## 14.2 Message queue operations

### 14.2.1 MESSAGE\_DELETE

```
MESSAGE_DELETE (
    queue      : Message_queue_designator,
    position    : Natural
)
```

MESSAGE\_DELETE removes the message with message number *position* from the message queue *queue*. The space used of *queue* is decremented by the space used by the removed message, and the message count of *queue* is decremented by 1.

A read lock of the default mode is obtained on *queue*.

A read lock of the default mode is obtained on *queue*.

#### Errors

ACCESS\_ERRORS (*queue*, ATOMIC, MODIFY, WRITE, CONTENTS)

MESSAGE\_POSITION\_IS\_NOT\_VALID (*position*, *queue*)

MESSAGE\_QUEUE\_IS\_RESERVED (*queue*)

### 14.2.2 MESSAGE\_PEEK

```
MESSAGE_PEEK (
    queue      : Message_queue_designator,
    types      : Message_types,
    position    : [ Natural ]
)
    message : [ Received_message ]
```

MESSAGE\_PEEK reads a message from the message queue *queue* without removing it from *queue*. *types* specifies the set of acceptable message types. *position* specifies a position in *queue*; if it is 0 or not supplied, the position is the beginning of *queue*; otherwise it is the position immediately before the message with position number *position*.

If *queue* contains no messages of an acceptable message type after the specified position, then no message is returned. Otherwise a copy of the next message of an acceptable message type after the specified position is returned. In either case *queue* is unchanged.

#### Errors

ACCESS\_ERRORS (*queue*, ATOMIC, READ, READ\_CONTENTS)

MESSAGE\_POSITION\_IS\_NOT\_VALID (*position*, *queue*)

MESSAGE\_QUEUE\_IS\_RESERVED (*queue*)

### 14.2.3 MESSAGE\_RECEIVE\_NO\_WAIT

```
MESSAGE_RECEIVE_NO_WAIT (
    queue      : Message_queue_designator,
    types      : Message_types,
    position    : [ Natural ]
)
    message : [ Received_message ]
```

MESSAGE\_RECEIVE\_NO\_WAIT reads and removes a message from the message queue *queue*, but does not wait if there is no message of an acceptable message type in *queue*. *types* specifies the set of acceptable message types. *position* specifies a position in *queue*; if it is 0 or not supplied, the position is the beginning of the *queue*; otherwise it is the position immediately before the message with position number *position*.

If the *queue* contains no messages of an acceptable message type after the specified position, then no message is returned. Otherwise the first message of an acceptable type after the specified position is returned, and that message is removed from *queue*. The last receive time of *queue* is set to the system time, the space used of *queue* is decremented by the space used of the removed message, and the message count of *queue* is decremented by 1.

A read lock of the default mode is obtained on *queue*.

### Errors

ACCESS\_ERRORS (*queue*, ATOMIC, MODIFY, WRITE\_CONTENTS)

MESSAGE\_POSITION\_IS\_NOT\_VALID (*position*, *queue*)

MESSAGE\_QUEUE\_IS\_RESERVED (*queue*)

MESSAGE\_TYPES\_NOT\_FOUND\_IN\_QUEUE (*queue*, *types*, *position*)

## 14.2.4 MESSAGE\_RECEIVE\_WAIT

```
MESSAGE_RECEIVE_WAIT (
    queue      : Message_queue_designator,
    types      : Message_types,
    position   : [ Natural ]
)
    message   : Received_message
```

MESSAGE\_RECEIVE\_WAIT reads and removes a message from the message queue *queue*, waiting if necessary for a message of an acceptable message type to arrive. *types* specifies the set of acceptable message types. *position* specifies a position in the message queue *queue*; if it is 0 or not supplied, the position is the beginning of the *queue*; otherwise it is the position immediately before the message with position number *position*.

If the message queue *queue* contains one or more messages of an acceptable message type after the specified position, the first such message is returned, and that message is removed from the queue. The last receive time of *queue* is set to the system time, the space used of *queue* is decremented by the space used of the removed message, and the message count of *queue* is decremented by 1.

If *queue* contains no messages of any acceptable message type after the specified position, then the operation waits and reader waiting for *queue* is set to **true**, until one of the following happens.

- A message of an acceptable type is placed on the queue. If the calling process is the listening process for *queue* and the message type of the message is one of the specified set of message types for the associated handler, then the handler is executed; otherwise the operation proceeds as described above.
- A reserved message queue of the calling process receives a message of message type WAKE. The error condition MESSAGE\_QUEUE\_HAS\_BEEN\_WOKEN then holds.

- The message queue is removed from the object base. The error condition MESSAGE\_QUEUE\_HAS\_BEEN\_DELETED then holds.
- The caller is denied mandatory read access, mandatory write access, or WRITE\_CONTENTS discretionary access to *queue*. The error condition CONFIDENTIALITY\_WOULD\_BE\_VIOLATED or INTEGRITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED then holds in the first case, CONFIDENTIALITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED or INTEGRITY\_WOULD\_BE\_VIOLATED in the second case, and DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED in the third case (all under ACCESS\_ERRORS).
- The message queue becomes reserved by another process. The error condition MESSAGE\_QUEUE\_IS\_RESERVED then holds.

A read lock of the default mode is obtained on *queue*.

#### Errors

ACCESS\_ERRORS (*queue*, ATOMIC, MODIFY, WRITE\_CONTENTS)

MESSAGE\_POSITION\_IS\_NOT\_VALID (*position*, *queue*)

MESSAGE\_QUEUE\_IS\_RESERVED (*queue*)

MESSAGE\_QUEUE\_HAS\_BEEN\_DELETED (*queue*)

MESSAGE\_QUEUE\_HAS\_BEEN\_WOKEN (*queue*)

### 14.2.5 MESSAGE\_SEND\_NO\_WAIT

```
MESSAGE_SEND_NO_WAIT (
    queue      : Message_queue_designator,
    message    : Message
)
```

MESSAGE\_SEND\_NO\_WAIT appends the message *message* to the message queue *queue*.

The last send time of *queue* is set to the system time. The space used of *queue* is incremented by the space used by *message*. The message count of *queue* is incremented by 1.

A read lock of the default mode is obtained on *queue*.

#### Errors

ACCESS\_ERRORS (*queue*, ATOMIC, MODIFY, APPEND\_CONTENTS)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_MESSAGE\_SIZE)

MESSAGE\_QUEUE\_WOULD\_BE\_TOO\_BIG (*queue*)

### 14.2.6 MESSAGE\_SEND\_WAIT

```
MESSAGE_SEND_WAIT (
    queue      : Message_queue_designator,
    message    : Message
)
```

MESSAGE\_SEND\_WAIT appends the message *message* to the message queue *queue*, waiting if necessary until *queue* has enough space for it.

If the space used of the message queue *queue* would not exceed the total space of *queue*, the message *message* is appended to *queue*. The last send time of *queue* is set to the system time when the message is sent (at the end of the waiting period, if any). The space used of *queue* is incremented by the space used by *message*. The message count of *queue* is incremented by 1.

If the space used of *queue* would exceed the total space of *queue*, then writer waiting of *queue* is set to **true** and the operation waits until one of the following occurs.

- The space used of *queue* would no longer exceed the total space of *queue*. The operation then proceeds as described above.
- The calling process receives a message. of message type WAKE. The error condition MESSAGE\_QUEUE\_HAS\_BEEN\_WOKEN then holds.
- The message queue is removed from the object base. The error condition MESSAGE\_QUEUE\_HAS\_BEEN\_DELETED then holds.
- The caller is denied mandatory read access, mandatory write access, or APPEND\_CONTENTS discretionary access. The error condition CONFIDENTIALITY\_WOULD\_BE\_VIOLATED or INTEGRITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED then holds in the first case, CONFIDENTIALITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED or INTEGRITY\_WOULD\_BE\_VIOLATED in the second case, and DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED in the third case (all under ACCESS\_ERRORS).

A read lock of the default mode is obtained on *queue*.

#### Errors

ACCESS\_ERRORS (*queue*, ATOMIC, MODIFY, APPEND\_CONTENTS)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_MESSAGE\_SIZE)

MESSAGE\_QUEUE\_HAS\_BEEN\_DELETED (*queue*)

MESSAGE\_QUEUE\_HAS\_BEEN\_WOKEN (*queue*)

### 14.2.7 QUEUE\_EMPTY

```
QUEUE_EMPTY (
    queue : Message_queue_designator
)
```

QUEUE\_EMPTY empties the message queue *queue*, i.e. removes all messages from it.

A read lock of the default mode is obtained on *queue*.

#### Errors

ACCESS\_ERRORS (*queue*, ATOMIC, MODIFY, WRITE\_CONTENTS)

MESSAGE\_QUEUE\_IS\_RESERVED (*queue*)

### 14.2.8 QUEUE\_HANDLER\_DISABLE

```
QUEUE_HANDLER_DISABLE (
    queue : Message_queue_designator
)
```

QUEUE\_HANDLER\_DISABLE makes the calling process no longer the listening process for the message queue *queue*. *queue* must be reserved by the calling process.

The "is\_listener" link from the calling process to *queue* and its reverse are deleted.

#### Errors

ACCESS\_ERRORS (*queue*, ATOMIC, MODIFY, WRITE\_LINKS)

MESSAGE\_QUEUE\_HAS\_NO\_HANDLER (*queue*)

MESSAGE\_QUEUE\_IS\_NOT\_RESERVED (*queue*)

### 14.2.9 QUEUE\_HANDLER\_ENABLE

```
QUEUE_HANDLER_ENABLE (
    queue      : Message_queue_designator,
    types      : Message_types,
    handler    : Handler
)
```

QUEUE\_HANDLER\_ENABLE makes the calling process the listening process for the message queue *queue*, with associated message types specified by *types*, and handler *handler*.

An "is\_listener" link is created from the calling process to *queue*, with "message\_types" attribute set to a value representing *types*.

The previous handler, if any, for *queue* is disabled as by a prior call of QUEUE\_HANDLER\_DISABLE.

#### Errors

ACCESS\_ERRORS (*queue*, ATOMIC, MODIFY, WRITE\_LINKS)

MESSAGE\_QUEUE\_IS\_NOT\_RESERVED (*queue*)

### 14.2.10 QUEUE\_RESERVE

```
QUEUE_RESERVE (
    queue : Message_queue_designator
)
```

QUEUE\_RESERVE reserves the message queue *queue* for the calling process. If *queue* is already reserved for the calling process, QUEUE\_RESERVE has no effect.

A "reserved\_message\_queue" link reversed by a "reserved\_by" link is created from the calling process to *queue*.

#### Errors

ACCESS\_ERRORS (*queue*, ATOMIC, MODIFY, APPEND\_LINKS)

MESSAGE\_QUEUE\_IS\_RESERVED (*queue*)

### 14.2.11 QUEUE\_RESTORE

```

QUEUE_RESTORE (
    queue : Message_queue_designator,
    file   : File_designator
)

```

QUEUE\_RESTORE reconstructs the message queue *queue* from the contents of the object designated by *file*.

The last access time of *file* is set to the system time.

A write lock of the default mode is obtained on *queue* and a read lock of the default mode is obtained on *file*.

#### Errors

ACCESS\_ERRORS (*file*, ATOMIC, READ, READ\_CONTENTS)  
 ACCESS\_ERRORS (*queue*, ATOMIC, MODIFY, WRITE\_CONTENTS)  
 CONTENTS\_FORMAT\_IS\_INVALID (*file*)  
 LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_MESSAGE\_QUEUE\_SPACE)  
 MESSAGE\_QUEUE\_IS\_BUSY (*queue*)  
 MESSAGE\_QUEUE\_IS\_RESERVED (*queue*)  
 MESSAGE\_QUEUE\_WOULD\_BE\_TOO\_BIG (*queue*)

### 14.2.12 QUEUE\_SAVE

```

QUEUE_SAVE (
    queue : Message_queue_designator,
    file   : File_designator
)

```

QUEUE\_SAVE copies all the messages from the message queue *queue* to the contents of the file *file*. The existing contents of *file* is overwritten. The format of the contents of *file* is implementation-defined.

The message queue *queue* is unaffected, except that any "notifier" links from *queue* are deleted.

The last change time and last modification time of *file* are set to the system time of the call.

A write lock of the default mode is obtained on *file* and a read lock on *queue*.

#### Errors

ACCESS\_ERRORS (*file*, ATOMIC, MODIFY, WRITE\_CONTENTS)  
 ACCESS\_ERRORS (*queue*, ATOMIC, READ, READ\_CONTENTS)  
 MESSAGE\_QUEUE\_IS\_RESERVED (*queue*)

### 14.2.13 QUEUE\_SET\_TOTAL\_SPACE

```

QUEUE_SET_TOTAL_SPACE (
    queue       : Message_queue_designator,
    total_space : Natural
)

```



QUEUE\_SET\_TOTAL\_SPACE sets the total space of the message queue *queue* to the value of *total\_space*.

A write lock of the default mode is obtained on *queue*.

#### Errors

ACCESS\_ERRORS (*queue*, ATOMIC, CHANGE, CONTROL\_OBJECT)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_MESSAGE\_QUEUE\_SPACE)

MESSAGE\_QUEUE\_IS\_RESERVED (*queue*)

MESSAGE\_QUEUE\_TOTAL\_SPACE\_WOULD\_BE\_TOO\_SMALL (*queue*, *total\_space*)

### 14.2.14 QUEUE\_UNRESERVE

```
QUEUE_UNRESERVE (
    queue : Message_queue_designator
)
```

QUEUE\_UNRESERVE unreserves the message queue *queue* for the calling process. If *queue* is not reserved for the calling process, QUEUE\_UNRESERVE has no effect.

The "reserved\_message\_queue" and "reserved\_by" links between the calling process and *queue* are deleted. If the calling process has an "is\_listener" link to *queue* then that link and its reverse "listened\_to" link are deleted.

#### Errors

ACCESS\_ERRORS (*queue*, ATOMIC, MODIFY, WRITE\_LINKS)

NOTE - The termination of a process implies the unreserving of all the process's reserved message queues.

## 15 Notification

### 15.1 Notification concepts

#### 15.1.1 Access events and notifiers

Access\_event = MODIFICATION\_EVENT | CHANGE\_EVENT | DELETE\_EVENT | MOVE\_EVENT

Access\_events = **set of** Access\_event

A *notifier* is a "notifier" link from a message queue to an object with key attribute "notifier\_key" and the attributes "modification\_event", "change\_event", "delete\_event", and "move\_event", referred to as *monitored access attributes*. See 14.1 for the DDL definition of notifiers.

A *monitored object* is a destination object of a notifier. The values of the monitored access attributes of the notifier define the events on which the monitored object is monitored:

- Modification event is **true**. Modification events: an operation implicitly sets the last modification time of the object.
- Change event is **true**. Change events: an operation implicitly sets the last change time but not the last modification time of the object. If the operation only sets the volume identifier of the object, the CHANGE\_EVENT event is not raised.
- Delete event is **true**. Delete events: an operation results in the deletion of the object.

- Move event is **true**. Move events: an operation results in a change to the volume identifier of the object, including archiving the object and restoring it from archive.

The notification mechanism sends notification messages to message queues when a specified access is carried out on a monitored object. The specified access event is said to be *raised* by the operation that accessed the object. The notification mechanism is said to be *triggered* by the raised event.

If there is a notifier from a message queue to any object then that message queue has a reserving process.

#### NOTES

- 1 The monitored access attributes of the notifier define the access events for which the destination of the notifier is to be monitored. Their initial value is **false**. For each attribute, if the value of the attribute is **true**, then the object is being monitored for that event.
- 2 Each value of the notifier key identifies a specific notifier in the context of the associated message queue. As implied by the DDL specification, the notifier key is unique in the context of the associated message queue.
- 3 In order to carry out notification mechanism operations, a process must reserve the message queue which is to be used as recipient of the notification messages, and in order to be notified, the message queue must remain reserved by the process.
- 4 If a process unreserves a message queue then any notifiers from the message queue are deleted.
- 5 A process can reserve several different message queues for notification purposes. For each of these message queues, it can create several notifiers (one for each object under monitoring). An object can be monitored using several message queues by one process or by several processes.
- 6 There are additional possibilities for the deletion and moving of objects other than by the OBJECT\_DELETE and OBJECT\_MOVE operations.

### 15.1.2 Notification messages

Notification\_message\_type = MODIFICATION\_MSG | CHANGE\_MSG | DELETE\_MSG |  
MOVE\_MSG | NOT\_ACCESSIBLE\_MSG | LOST\_MSG

A *notification message* is a message (see 14.1) sent by the notification mechanism to one or more message queues each time an object under monitoring is accessed in a way which has been specified to be monitored. The type of such a message specifies the access event that has been carried out on the monitored object or the information that the monitored object is no longer accessible or that modification messages have been lost. The possible values of the type of a notification message are defined as follows:

- MODIFICATION\_MSG: Notifies that a modification access event has been raised (except CONTENTS\_WRITE or CONTENTS\_TRUNCATE).
- CHANGE\_MSG: Notifies that a change access event has been raised.
- DELETE\_MSG: Notifies that a delete access event has been raised.
- MOVE\_MSG: Notifies that a move access event has been raised.
- NOT\_ACCESSIBLE\_MSG: A message of this type is sent to a message queue each time a monitored object becomes no longer accessible from the workstation on which that message queue resides.

- **LOST\_MSG:** When a message queue is full and there is not sufficient space on the queue to store a notification message, the notification messages to be sent by the notification mechanism are lost. In this case, when the message queue empties sufficiently to give space for a notification message, a message is sent to the message queue by the notification mechanism saying that some messages have been lost.

The data of the message includes in an implementation-defined way the notifier key that associates the message queue and the monitored object.

At most four notification messages are sent to a message queue, one for each type of access carried out on the object during the period it was explicitly locked. The order of these four messages is implementation-defined.

When a monitored object is archived, a message of type **MOVE\_MSG** and a message of type **NOT\_ACCESSIBLE\_MSG** are both sent; when a monitored object is restored from archive, a message of type **MOVE\_MSG** is sent.

### 15.1.3 Time of sending notification messages

The end of an operation and the releasing of a lock define the points in time at which the notification messages are sent to processes as defined in 15.1.4.

At the appropriate point in time, the switched on access events are raised, triggering the notification mechanism which sends the notification messages to the message queues associated by notifiers with the object that has had been modified, changed, deleted, or moved.

A message of message type **NOT\_ACCESSIBLE\_MSG** is sent by the notification mechanism when **WORKSTATION\_REDUCE\_CONNECTION** or **WORKSTATION\_DISCONNECT** is called or when a network partition is detected, such that the monitored object becomes inaccessible in the specified manner.

### 15.1.4 Range of concerned message queues

For an operation modifying, changing, deleting, or moving an object, a notification message is sent to all the message queues associated with that object by a notifier when the update becomes available to the process reserving the message queue.

If the message queue security labels are such that writing to the queue by the process accessing the object would give rise to a mandatory security violation, then no notification message is sent.

On transaction rollback, notification messages are sent notifying rollback and no messages are sent to non-enclosed activities.

## 15.2 Notification operations

### 15.2.1 NOTIFICATION\_MESSAGE\_GET\_KEY

```

NOTIFICATION_MESSAGE_GET_KEY(
    message      : Message,
)
    notifier_key : Natural

```

NOTIFICATION\_MESSAGE\_GET\_KEY returns a notifier key *notifier\_key* derived from the data of the notification message *message*.

*notifier\_key* is the notifier key of the notifier whose monitored object underwent the access event which triggered the sending of *message*.

### Errors

MESSAGE\_IS\_NOT\_A\_NOTIFICATION\_MESSAGE (*message*)

NOTE - The notifier identified by *notifier\_key* may no longer exist.

## 15.2.2 NOTIFY\_CREATE

```
NOTIFY_CREATE (
    notifier_key : Natural,
    queue       : Message_queue_designator,
    object      : Object_designator
)
```

NOTIFY\_CREATE creates a notifier from the message queue *queue* to the object *object*.

The notifier key of the notifier is set to *notifier\_key*. The monitored access attributes of the notifier are all set to **false**.

### Errors

ACCESS\_ERRORS (*queue*, ATOMIC, MODIFY, APPEND\_LINKS)

CONFIDENTIALITY\_WOULD\_BE\_VIOLATED (*object*, ATOMIC)

INTEGRITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED (*object*, ATOMIC)

MESSAGE\_QUEUE\_IS\_NOT\_RESERVED (*queue*)

NOTIFIER\_KEY\_EXISTS (*notifier\_key*)

OBJECT\_IS\_INACCESSIBLE (*object*, ATOMIC)

OBJECT\_IS\_ARCHIVED (*object*)

### NOTES

- 1 The creation of a notifier from a message queue to an object means that a notification message will be sent to the message queue whenever the object is accessed with some specified access events.
- 2 Initially, on creation, the monitored access attributes are set to **false**, so no events are specified. The monitored access events may be changed by NOTIFY\_SWITCH\_EVENTS.
- 3 As implied by the DDL specification, the *notifier\_key* value must be unique in the context of the message queue and must be greater than or equal to zero; apart from these constraints, it may be freely chosen by the user.

## 15.2.3 NOTIFY\_DELETE

```
NOTIFY_DELETE (
    notifier_key : Natural,
    queue       : Message_queue_designator
)
```

NOTIFY\_DELETE deletes the notifier with notifier key *notifier\_key* from the message queue *queue*.

**Errors**ACCESS\_ERRORS (*queue*, ATOMIC, MODIFY, WRITE\_LINKS)MESSAGE\_QUEUE\_IS\_NOT\_RESERVED (*queue*)NOTIFIER\_KEY\_DOES\_NOT\_EXIST (*notifier\_key*)

NOTE - The object which was monitored by *notifier* may continue to be monitored by other notifiers into other message queues. Other objects may continue to be monitored by other notifiers associated with *queue*.

**15.2.4 NOTIFY\_SWITCH\_EVENTS**

```

NOTIFY_SWITCH_EVENTS (
    notifier_key    : Natural,
    queue           : Message_queue_designator,
    access_events : Access_events
)

```

NOTIFY\_SWITCH\_EVENTS sets each of the monitored access attributes of the notifier with notifier key *notifier\_key* from the message queue *queue* to **true** if the corresponding access event is in *access\_events*, and to **false** otherwise.

**Errors**ACCESS\_ERRORS (*queue*, ATOMIC, MODIFY, WRITE\_LINKS)MESSAGE\_QUEUE\_IS\_NOT\_RESERVED (*queue*)NOTIFIER\_KEY\_DOES\_NOT\_EXIST (*notifier\_key*)

NOTE - Switching on an access event of a notifier (setting the attribute value to true) means that the associated object is then under monitoring for that access event. Switching off an access event (setting the attribute value to false) means that the associated object is no longer under monitoring for that access event.

**16 Concurrency and integrity control****16.1 Concurrency and integrity control concepts****16.1.1 Activities**

Activity\_class = UNPROTECTED | PROTECTED | TRANSACTION

sds system:

activity\_class: (read) **enumeration** (UNPROTECTED, PROTECTED, TRANSACTION) := UNPROTECTED;

activity\_status: (read) **non\_duplicated enumeration** (UNKNOWN, ACTIVE, COMMITTING, ABORTING, COMMITTED, ABORTED) := UNKNOWN;

activity: **child type of object with attribute**

activity\_class;  
 activity\_status;  
 activity\_start\_time: **(read) time**;  
 activity\_termination\_start\_time: **(read) time**;  
 activity\_termination\_end\_time: **(read) time**;

**link**

started\_by: **(navigate) reference link to process reverse** started\_activity;  
 nested\_in: **(navigate) reference link to activity reverse** nested\_activity;  
 nested\_activity: **(navigate) implicit link (system\_key) to activity reverse** nested\_in;  
 process\_started\_in: **(navigate) implicit link (system\_key) to process reverse** started\_in\_activity;

**end** activity;

**end** system;

An activity is the framework in which a set of related operations takes place. Each operation is always carried out on behalf of just one activity. An activity is started at the time it is created and remains in existence until the deactivation of the process which started it.

The activity class of an activity describes the degree of protection which the activity requires; it affects the default level of concurrency control applicable to operations carried out on behalf of the activity. There are three activity classes:

- UNPROTECTED. An *unprotected* activity, used when it is not necessary to protect data from concurrent activities.
- PROTECTED. A *protected* activity, used when data to be accessed needs protection from concurrent activities.
- TRANSACTION. A *transaction* activity (or *transaction*), used when the activity has a significant effect on the object base and its integrity needs to be protected.

The activity status records the current state of the activity. The possible states of an activity are:

- UNKNOWN. The "activity" object has been created by an operation defined in clause 9.
- ACTIVE. The activity is started and its termination is not yet initiated.
- COMMITTING. The activity's normal termination is initiated but not completed.
- ABORTING. The activity's abnormal termination is initiated but not completed.
- COMMITTED. The activity is normally terminated.
- ABORTED. The activity is abnormally terminated.

The activity start time records the time when the activity was started.

The activity termination start time records the time when the termination of the activity was started.

The activity termination end time records the time when the termination of the activity was completed.

The "started\_by" process is the process that started the activity.

The "nested\_in" activity, called the *enclosing activity* of the activity, is the activity within which the activity was started. The *nested activities* of an activity are the activities for which the activity is the enclosing activity.

The "process\_started\_in" processes are the processes which were created while the activity was the current activity.

Within each process there is only one *current activity*. When a process is initiated, its current activity is the current activity of its parent process. When an activity is started in a process it becomes the current activity of the process; the current activity is then the activity of the process with the highest key in the "started\_activity" link from the process and which is still active. When an activity is terminated in a process its immediate enclosing activity becomes the current activity of the process.

Each workstation in a PCTE installation has an outermost activity. The *outermost activity* of a workstation is an unprotected activity that is created by implementation-dependent means such that it is indistinguishable from an activity created by ACTIVITY\_START except that it has no "started\_by" or "nested\_in" link. It has a "process\_started\_in" link to the initial process.

Updates by an activity to a resource are *available* if, when the updated resource is read by another activity not enclosed by the updating activity, data derived from the updated state of the resource is obtained. Data derived from the updated state of the resource is obtained when read by the updating activity and by nested activities without necessarily being generally available.

#### NOTES

- 1 Activities can be internal to one process or can extend over several descendant processes. A process is free to start an activity, but a process is only allowed to terminate activities that it has started.
- 2 Operations performed by a process, other than those on an open contents, are carried out on behalf of the current activity of the process at the time the operation is called.
- 3 A nested transaction may be terminated without implying the termination of its enclosing transaction. When a transaction is normally terminated then all the read locks it has acquired are released and all the write locks of default mode it has acquired or inherited from its nested transactions are inherited by its enclosing transaction. When a transaction is abnormally terminated then the changes made by it and all its nested transactions are unmade (unless explicitly excluded from rollback) and all the locks it has acquired, including the write locks it has inherited from its nested transactions, are released. This effect is transitive so that, in the case of successive normal terminations of transactions nested one in another, nested transaction write locks are not released, and the changes not committed, until the outermost transaction is normally terminated.
- 4 Protected or unprotected activities may also be nested within transactions. In this case, modifications made within the nested activities are considered also to be changes made within their closest enclosing transaction. Accordingly, when locks are acquired by nested protected or unprotected activities, locks are implicitly acquired at the same time by their closest enclosing transaction (see 16.1.6)
- 5 In the same way when a lock whose mode is not the default write mode is acquired by a nested transaction, a lock is implicitly acquired at the same time by the closest enclosing transaction.
- 6 A process running on behalf of a transaction can explicitly exclude from rollback changes made to certain resources by explicitly locking such resources in unprotected or protected modes (i.e. not in default write modes) (see 16.1.5). However creating or deleting of objects and links cannot be excluded from rollback.
- 7 The outermost activity of a workstation is implicitly set up by the system. It is intended to provide a valid activity framework for the initial process of the workstation. Each workstation has its own outermost activity, i.e. the outermost activity of a workstation cannot be the outermost activity of another workstation. An initial process is initiated in the context of that activity. It is intended that the initial process should then start an activity suitable for its own requirements.



8 Transactions do not protect the local data of a process, hence, for example, changes to contents handles, object references, and other local variables made within the scope of a transaction are not reversed if the transaction is aborted.

### 16.1.2 Resources and locks

```
Lock_internal_mode = READ_UNPROTECTED | READ_SEMIPROTECTED |
WRITE_UNPROTECTED | WRITE_SEMIPROTECTED | DELETE_UNPROTECTED |
DELETE_SEMIPROTECTED | READ_PROTECTED | DELETE_PROTECTED |
WRITE_PROTECTED
```

```
Lock_set_mode = Lock_internal_mode | WRITE_TRANSACTIONED |
DELETE_TRANSACTIONED | READ_DEFAULT | WRITE_DEFAULT | DELETE_DEFAULT
```

**sds** system:

```
lock_mode: READ_UNPROTECTED, READ_SEMIPROTECTED, WRITE_UNPROTECTED,
WRITE_SEMIPROTECTED, DELETE_UNPROTECTED, DELETE_SEMIPROTECTED,
READ_PROTECTED, DELETE_PROTECTED, WRITE_PROTECTED,
WRITE_TRANSACTIONED, DELETE_TRANSACTIONED;
```

```
lock_external_mode: (read) enumeration (lock_mode) := READ_UNPROTECTED;
```

```
lock_internal_mode: (read) enumeration (lock_external_mode range
READ_UNPROTECTED .. WRITE_PROTECTED) := READ_UNPROTECTED;
```

**extend object type** activity with

**link**

lock: (navigate) non\_duplicated designation link (lock\_identifier) to object with attribute

locked\_link\_name;

lock\_external\_mode;

lock\_internal\_mode;

lock\_explicitness: (read) enumeration (EXPLICIT, IMPLICIT) := IMPLICIT;

lock\_duration: (read) enumeration (SHORT, LONG) := SHORT;

**end** lock;

**end** activity;

**extend link type** process\_waiting\_for with

**attribute**

lock\_external\_mode;

lock\_internal\_mode;

**end** process\_waiting\_for;

**end** system;

A *resource* is either an object resource or a link resource.

An *object resource* is an object restricted to the following:

- its contents,
- its type,
- its preferred link type and preferred link key,
- its attributes, except the predefined attributes "last\_access\_time", "last\_change\_time", "last\_modification\_time", "last\_composite\_access\_time", "last\_composite\_change\_time", "last\_composite\_modification\_time", "num\_incoming\_links", "num\_incoming\_composition\_links", "num\_incoming\_existence\_links", "num\_incoming\_reference\_links",

- "num\_incoming\_stabilizing\_links", "num\_outgoing\_composition\_links", and  
 "num\_outgoing\_existence\_links",
- its incoming "object\_on\_volume" link.

A *link resource* is a link, identified by its link name and restricted to the following:

- its link type,
- its sequence of key attributes,
- its set of non-key attributes,
- the object designator of its destination.

The fact that a resource is locked is represented by a "locked\_by" link from the object resource or the origin of a link resource to the activity which holds the lock. The locked link name of the link specifies whether the resource is an object resource or a link resource:

- if the locked resource is a link resource, the "locked\_link\_name" attribute is set to the link name in canonical form (see 23.1.2.4).
- if the locked resource is an object resource, the "locked\_link\_name" attribute is set to the empty string.

A *lock* is represented by a "lock" link from an activity to a resource; the activity is said to *hold* the lock *on* the resource. The link is created at the time the lock is established and remains until the lock is released or inherited. Locks ensure the consistency of object base data access operations by controlling the synchronization of concurrent operations on the same resources.

A lock is characterized by a unique lock identifier, the value of which is implementation-dependent.

The *concerned domain* of a resource is the set of resources which can be affected by modifications of that resource:

- if the resource is an object, the concerned domain is the object resource and the set of links (link resources) originating from the object.
- if the resource is a link, the concerned domain is the link resource and the object (object resource) from which the link starts.

A resource is said to be *operated on* by an activity when:

- either the resource is an object whose contents are currently open (see clause 12), by CONTENTS\_OPEN or PROCESS\_START on behalf of that activity, in which case the resource is operated on while the contents is open;
- or the resource (i.e. object or link) is the subject of operations other than operations on "lock" and "locked\_by" links and on the contents of objects, in which case the resource is operated on for the duration of the operation.

An activity can lock a resource just once; i.e. two locks originating from the same activity cannot have the same locked resource and the same destination.

A lock has the following attributes:

- A lock external mode, which controls synchronization of resource accesses between an activity and all other activities which are not nested (either directly or transitively) to it.

- A lock internal mode, which controls synchronization of resource accesses between an activity and all activities which are nested (either directly or transitively) to it.

The lock internal mode is equal to or weaker than the lock external mode (see below). See below for a definition of lock modes.

- A lock explicitness, which records how the lock was established:
  - . EXPLICIT. An *explicit* lock, i.e. it was established explicitly by one of locking operations.
  - . IMPLICIT. An *implicit* lock, i.e. it was established implicitly as the resource was implicitly acquired.
- A lock duration, which records the duration of the lock:
  - . LONG. A *long* lock, i.e. one which, once established, holds until the termination of the activity.
  - . SHORT. A *short* lock, i.e. one which can be released before the termination of the activity.

A long lock can be held only by a transaction.

A short lock can be held only by a protected or an unprotected activity.

NOTE – The incoming "object\_on\_volume" links of an object resource are created or deleted by operations which create, move, or delete objects. When these operations are performed in a transaction which is then rolled back, the creation or deletion of these links is also rolled back.

### 16.1.3 Lock modes

The meanings of the lock mode values are as follows. The abbreviations shown are used in the tables at the end of this clause.

- READ\_UNPROTECTED (RUN). The activity holding the lock can read the resource. Other activities can concurrently read or write to the same resource or delete it.
- READ\_SEMIPROTECTED (RSP) (for object resources only). The activity holding the lock can read the resource. Other activities can concurrently read from the same resource. Other activities can concurrently read or write to the same resource with WRITE\_UNPROTECTED, WRITE\_SEMIPROTECTED, WRITE\_PROTECTED or WRITE\_TRANSACTIONED locks but cannot delete it.
- WRITE\_UNPROTECTED (WUN). The activity holding the lock can read or write to the resource.

If the resource is an object, other activities can concurrently read or write to the same resource or delete it with DELETE\_UNPROTECTED or DELETE\_SEMIPROTECTED locks, other activities can concurrently read or write to the same resource with WRITE\_UNPROTECTED or WRITE\_SEMIPROTECTED locks and other activities can concurrently read the resource with READ\_UNPROTECTED or READ\_SEMIPROTECTED locks.

If the resource is a link, other activities can concurrently read or write to the same resource or delete it with WRITE\_UNPROTECTED locks and other activities can concurrently read the resource with READ\_UNPROTECTED locks.

Updates to the resource are available if an enclosing activity does not hold a WTR or DTR lock on the resource.

- WRITE\_SEMIPROTECTED (WSP) (for object resources only). The activity holding the lock can read or write to the resource. Other activities can concurrently read or write to the same resource with WRITE\_UNPROTECTED or WRITE\_SEMIPROTECTED locks but cannot delete it and other activities can concurrently read the resource with READ\_UNPROTECTED or READ\_SEMIPROTECTED locks. Updates to the resource are available if an enclosing activity does not hold a WTR or DTR lock on the resource.
- DELETE\_UNPROTECTED (DUN) (for object resources only). The activity holding the lock can read or write to the resource or delete it. Other activities can concurrently read or write to the same resource or delete it with DELETE\_UNPROTECTED locks, other activities can concurrently read or write to the same resource with WRITE\_UNPROTECTED and other activities can concurrently read the resource with READ\_UNPROTECTED locks. Updates to the resource are available if an enclosing activity does not hold a WTR or DTR lock on the resource.
- DELETE\_SEMIPROTECTED (DSP) (for object resources only). The activity holding the lock can read or write to the resource or delete it. Other activities can concurrently read or write to the same resource with WRITE\_UNPROTECTED locks but cannot delete it and other activities can concurrently read the resource with READ\_UNPROTECTED locks. Updates to the resource are available if an enclosing activity does not hold a WTR or DTR lock on the resource.
- READ\_PROTECTED (RPR). The activity holding the lock can read the resource. Other activities can concurrently read the same resource with READ\_UNPROTECTED, READ\_SEMIPROTECTED or READ\_PROTECTED locks. No other activities can concurrently write to the same resource.
- WRITE\_PROTECTED (WPR). The activity holding the lock can read or write to the resource. Other activities can concurrently read the same resource with READ\_UNPROTECTED or READ\_SEMIPROTECTED locks. No other activities can concurrently write to the same resource. Updates to the resource are available if an enclosing activity does not hold a WTR or DTR lock on the resource.
- DELETE\_PROTECTED (DPR) (for object resources only). The activity holding the lock can read or write to the resource or delete it. Other activities can concurrently read the same resource with READ\_UNPROTECTED locks. No other activities can concurrently write to the same resource. Updates to the resource are available if an enclosing activity does not hold a WTR or DTR lock on the resource.
- WRITE\_TRANSACTIONED (WTR). Transaction holding the lock can read or write to the resource. Other activities can concurrently read the same resource with READ\_UNPROTECTED or READ\_SEMIPROTECTED locks. No other activities can concurrently write to the same resource.
- DELETE\_TRANSACTIONED (DTR) (for object resources only). Transaction holding the lock can read or write to the resource or delete it. Other activities can concurrently read the same resource with READ\_UNPROTECTED locks. No other activities can concurrently write to the same resource.

It is implementation-defined whether or not updates to a resource are available if the updates are performed while an activity holds a WTR or DTR lock on the resource.

Locks of the following modes, whether internal or external, can be held only on an object resource: READ\_SEMIPROTECTED, WRITE\_SEMIPROTECTED, DELETE\_SEMIPROTECTED, DELETE\_PROTECTED, DELETE\_TRANSACTIONED.

The modes of a lock on a given resource must be compatible with the modes of locks held by other activities on resources in the concerned domain of that resource. Its external mode must be compatible with the external mode of all locks held on the resources in the concerned domain by other activities which are not enclosing, nor nested to the issuing activity, and with the internal mode of all locks already established by the enclosing activities on the resources in the concerned domain. Its internal mode must be compatible with the external modes of all locks already established by the nested activities on the resources in the concerned domain.

The lock modes are grouped into two categories:

- *Read lock modes:* READ\_UNPROTECTED, READ\_SEMIPROTECTED, READ\_PROTECTED
- *Write lock modes:* WRITE\_PROTECTED, WRITE\_TRANSACTIONED, WRITE\_UNPROTECTED, WRITE\_SEMIPROTECTED, DELETE\_PROTECTED, DELETE\_UNPROTECTED, DELETE\_SEMIPROTECTED, DELETE\_TRANSACTIONED

There are three relations defined between lock modes: *relative strength*, *relative weakness*, and *compatibility*. The relative strength relation between lock modes is defined by table 4. The relative weakness relation is the inverse of the relative strength relation (i.e. L1 is weaker than L2 if and only if L2 is stronger than L1). The compatibility relation is defined by table 5.

Updates to an accounting log or an audit file, the "message\_count", "last\_send\_time" and "last\_receive\_time" attributes of a message queue, and the "last\_access\_time" attribute of an object, are never made on behalf of the current activity.

If the current activity is a transaction, it may be terminated in one of two ways:

- For updates performed on behalf of the current activity while transaction locks were established, the operation ACTIVITY\_END, which *commits* the transaction, results in those updates becoming permanent, providing the transaction locks are not inherited (see 16.1.4).
- The operation ACTIVITY\_ABORT, which *aborts* the transaction, causes the updates performed on behalf of the current activity while transaction locks were established to be undone, apart from updates applied to contents of "pipe", "message\_queue", and "device" objects.

#### 16.1.4 Inheritance of locks

Inheritance of locks occurs only between transactions nested one in the other. A transaction inherits write locks of default modes (i.e. locks of modes WTR or DTR) from its (immediate) nested transactions each time such a nested transaction terminates normally (i.e. when it commits).

When a transaction T1 terminates, the lock it holds on a resource X is inherited by the nearest enclosing transaction T of T1. If T already holds a lock on X, the lock is promoted according to the rules of implicit promotion (see 16.1.5).

NOTE - Updates to a resource are committed or cancelled when there cease to be any WTR or DTR locks on that resource. When an enclosing transaction inherits WTR and DTR locks, it also inherits the updates. Normally, when there is no further enclosing transaction, updates are committed when the current transaction activity ends. If however the enclosing transaction T1 holds a non-transaction lock on a resource updated under enclosed transaction T, then when T ends the transaction lock is not inherited and neither are the updates which are committed. An exception to this is when a transaction T2 enclosing T1 exists and has a WTR or DTR lock on the resource; in this case the updates are inherited by T2 when T ends and are not committed at that point.

### 16.1.5 Establishment and promotion of locks

A lock is *requested* on a resource on behalf of an activity if an attempt is made to create a lock on that resource on behalf of that activity.

A lock is *established* on a resource when a lock is requested on the resource on behalf of an activity and no lock has yet been acquired by the activity.

A lock is *explicitly established* by means of operation LOCK\_SET\_OBJECT. A lock is *implicitly established* if the resource is implicitly acquired by some operation (other than LOCK\_SET\_OBJECT) operating on the resource and carried out on behalf of the activity. Locks of mode RSP, WSP, and DSP can only be established explicitly.

The modes of a lock, once established, can evolve either implicitly, according to the way the resource is operated on, or explicitly by means of the lock set operations.

The following enumerates, for each activity class, the implicit lock modes which are requested depending on the access performed on the acquired resource. Locks on link resources are always implicit and therefore always adopt default modes.

- Default external modes:
  - for unprotected activities the external mode is RUN when reading a resource, WUN when creating or updating a resource or deleting a link resource, and DUN when deleting an object resource;
  - for protected activities the external mode is RPR when reading a resource, WPR when creating or updating a resource or deleting a link resource, and DPR when deleting an object resource;
  - for transaction activities the external mode is RPR when reading a resource, WTR when updating a resource or creating or deleting a link resource, and DTR when creating or deleting an object resource.
- Default internal modes: for every activity class, internal modes are WUN for resources being created or updated and link resources being deleted, DUN for object resources being deleted, and RUN in all other cases.

The lock mode actually acquired depends on whether the lock is established or promoted. If it is established then the lock mode acquired is the default lock mode. If a promotion occurs, see below. Tables 8 and 9 summarize the default external modes.



When a lock is requested on a resource on behalf of an activity and a lock has already been acquired by the activity, then the lock may be promoted. To *promote* a mode of a lock is to transform it to a stronger mode which is compatible (as for the establishment of a new lock) with other locks on resources in the concerned domain. See table 6 and 16.1.7.

*Implicit promotion* of either or both the internal and the external modes of a lock occurs when an operation performing a write access (e.g. OBJECT\_SET\_ATTRIBUTE or CONTENTS\_OPEN with an opening mode allowing write access) is applied to a resource already acquired by the activity with a lock whose modes allow only read access to that resource, or when an operation deleting an object (e.g. LINK\_DELETE or OBJECT\_DELETE) is applied to an object resource already acquired by the activity with a lock whose modes do not allow deletion of that object resource.

*Explicit promotion* of either the internal or the external lock mode occurs when the lock set operations are applied to a resource already locked (either explicitly or implicitly) by the activity on behalf of which the lock set operation is carried out. The new mode must obey the promotion rules for lock modes (see below).

Any explicit attempt to promote an explicit or an implicit external or internal mode M1, or implicit attempt to promote an implicit external or internal mode M1, to a mode M2 which has no relation of relative strength with M1 is converted into an attempt to promote M1 to the weakest mode which is stronger than both M1 and M2 (e.g. an attempt to promote a RPR mode to a WUN mode is implicitly converted into an attempt to promote the mode to WPR). Table 6 defines the implicit promotion of lock modes when the prior lock is explicit; table 7 defines the promotion of lock modes for the other cases.

In the operation definitions, the phrase 'a *mode* lock of the default mode is *obtained on object*' where *object* is an object and *mode* is 'read' or 'write', is used to mean that an attempt is made to establish an implicit lock on the object resource *object* of a mode given by table 8, depending on the class of the current activity and the default lock mode of 'read' or 'write', and if 'write', whether *object* is being created, updated or deleted.

Similarly, the phrase 'a *mode* lock of the default mode is obtained on *link*' where *link* is a link and *mode* is 'read' or 'write', is used to mean that an attempt is made to establish an implicit lock on the link resource *link* of a mode given by table 9, depending on the class of the current activity and the default lock mode of 'read' or 'write', and if 'write', whether *link* is being created, updated or deleted. If *link* is a 'lock' or 'locked\_by' link then no lock is established on it.

If no lock currently exists on the resource (*object* or *link*) for that activity, an attempt is made to establish the lock, otherwise implicit promotion is attempted.

### 16.1.6 Implied locks

Locks can be established or promoted on a resource as a result of establishing or promoting another lock.

When locks are acquired by nested activities, this implies that implicit locks are acquired at the same time by their closest enclosing transaction:

- The establishing of (or promotion to) a WUN, WSP, WPR external mode lock for an activity also implies an implicit establishment (or implicit promotion) of a lock of external mode WTR on the resource on behalf of the closest enclosing transaction.



- The establishing of (or promotion to) a DUN, DSP, DPR external mode lock for an activity also implies an implicit establishment (or implicit promotion) of a lock of external mode DTR on the resource on behalf of the closest enclosing transaction.
- The establishing of (or promotion to) a RUN, RSP, RPR, WTR, or DTR external mode lock for an activity also implies an implicit establishment (or implicit promotion) of a lock of external mode RPR on the resource on behalf of the closest enclosing transaction.

Establishing or promoting a lock on a link also implies the implicit establishment (or promotion) of a read lock of the default mode on the origin of that link for the current activity, if the link type of the link has an upper bound or a non-zero lower bound and the link is being deleted or created.

The establishing of (or promotion to) a write lock on the last composition or existence link leading to an object for the purpose of its deletion also implies the implicit establishment (or implicit promotion) of a write lock allowing deletion on this object for the current activity (i.e. a DUN, DSP, DPR, or DTR lock according to the class of the activity and the promotion rules). The establishing of (or promotion to) a write lock on a composition or existence link for the purpose of deletion of the link results in the implicit establishment (or implicit promotion) of a read lock on the destination of the link.

In all these cases the internal lock mode of the implied lock is RUN.

#### 16.1.7 Conditions for establishment or promotion of a lock

The following conditions must be satisfied to establish or promote a lock on a given resource:

- Access rights: the current activity of a process can explicitly or implicitly establish a lock on a resource if and only if the process has at least one discretionary access right to the resource if it is an object resource; or to its origin if it is a link resource.
- Lock mode compatibility: an activity can establish (or promote) a lock on a resource if
  - its external mode is compatible with the external mode of all locks held on the resources in the concerned domain by other activities which are not enclosing, nor nested to the issuing activity;
  - its external mode is compatible with the internal mode of all locks already established by the enclosing activities on the resources in the concerned domain;
  - its internal mode is compatible with the external modes of all locks already established by the nested activities on the resources in the concerned domain;
  - The implied lock (if any) must also be compatible with existing locks as defined above.

If the conditions do not hold, either the issuing operation waits, waiting for the resource to become available, or the request returns an error without delay.

When an operation waits as a result of attempting to lock a resource in a mode which is incompatible with the existing locks on that resource held by other discrete activities, the operation is said to be *waiting on the resource*. When an operation is waiting on a resource, a "process\_waiting\_for" link is created from the process of the waiting operation to the resource. The "waiting\_type" attribute of that link is set to WAITING\_FOR\_LOCK, the required external and internal modes of the lock set in the "lock\_external\_mode" and "lock\_internal\_mode"

attributes respectively of the link, and the "locked\_link\_name" attribute is set to the link name, in canonical form, of the resource on which the lock is to be established if the resource is a link, and to the empty string otherwise. The link is removed when the operation which is waiting on the resource is interrupted, or the resource is acquired.

If a lock is held on an object resource by an activity, then any attempt to establish a lock on any of its links by that activity has no effect unless the lock mode resulting from the request is stronger than or has no relation of relative strength to the external mode of the lock on the object.

If a lock is held on a link resource by an activity, then when a lock is established on its origin by that activity the lock on the link is discarded, unless the external lock mode on the link is stronger than or has no relation of relative strength to the external mode of the lock on the origin.

### 16.1.8 Releasing locks

A distinction is made between *discarding* a lock (to get rid of it) and *releasing* a lock. Releasing a lock implies discarding the lock for the current activity, and if the lock has a WTR or a DTR mode then the closest enclosing transaction inherits the modifications to the resource. If there is no such transaction then modifications are *committed* (i.e. modifications can no longer be discarded).

In any case, this results in the deletion of the "lock" link and of the "locked\_by" link associated with the released lock. In the case that the lock is inherited by the closest enclosing transaction, if the resource was not already locked on behalf of that transaction, new "lock" and "locked\_by" links are created between this activity and the locked resource, in order to represent the inherited lock.

Long locks are released at the end of the activity. In the case of short locks two cases can apply:

- The lock was explicitly established: it is released either at the end of the activity or at the explicit unlock of the resource, whichever occurs first.
- The lock was implicitly established: it is released as soon as the locked resource is no longer being operated on behalf of the activity holding the lock (for example when the last open contents handle to the object contents is closed by CONTENTS\_CLOSE).

When a lock on a resource is discarded, if one or more operations are waiting on the resource then an attempt is made to establish or promote a lock on that resource in the modes given by the attributes "lock\_external\_mode" and "lock\_internal\_mode" of a "process\_waiting\_for" link to that resource on behalf of the current activity of the process which is the origin of that link. If a lock can be established or promoted on behalf of one of those activities, then the corresponding "process\_waiting\_for" link is deleted. If a lock can be established or promoted on behalf of more than one such activity, it is not defined on behalf of which activity it is established or promoted. The resource on which the lock is established or promoted is the destination of the "process\_waiting\_for" link if the "locked\_link\_name" attribute of that link is empty, otherwise it is the link with that attribute as link name.

#### NOTES

- 1 The description of each of the operations defines the resources, if any, which are operated on by the operation.
- 2 Nested parallel activities should be achieved by using parallel processes.

3 The internal mode of a lock held by an activity affects only the activity and its nested activities.

#### 16.1.9 Permanence of updates

When an update, whether made while a lock is established or not, is made *permanent*, the resulting change to objects in the object base is such that if a volume failure, device failure, or network failure event occurs so as to render one or more of those objects inaccessible, the objects retain their updated state. Conversely, if an update is not made permanent and such a failure event occurs then the objects revert to a state which existed before the update. An update to a link is considered to be an update to its origin.

If an update is made to an object or link while only non-transaction locks are established on that object or link then the update is made permanent at the latest when the activity in which the update occurred is terminated.

Updates made to objects or links, while WTR or DTR locks are established on those objects or links, are made permanent atomically when no transaction locks remain after a transaction end.

Updates made which do not require a lock to be established on the object or link, for example some operations defined in clause 13 and updates to audit files and accounting logs, are made permanent at an implementation-defined time.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 13719-1:1998

## 16.1.10 Tables for locks

Table 4 - Relative strength of lock modes

| <i>mode1</i>                                                                                                                                                                                                                                      | <i>mode2</i> |     |     |     |     |     |     |     |     |     |     |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
|                                                                                                                                                                                                                                                   | RUN          | RSP | WUN | WSP | RPR | WPR | WTR | DUN | DSP | DPR | DTR |
| RUN                                                                                                                                                                                                                                               | =            | <   | <   | <   | <   | <   | <   | <   | <   | <   | <   |
| RSP                                                                                                                                                                                                                                               | >            | =   | -   | <   | <   | <   | <   | -   | <   | <   | <   |
| WUN                                                                                                                                                                                                                                               | >            | -   | =   | <   | -   | <   | <   | <   | <   | <   | <   |
| WSP                                                                                                                                                                                                                                               | >            | >   | >   | =   | -   | <   | <   | -   | <   | <   | <   |
| RPR                                                                                                                                                                                                                                               | >            | >   | -   | -   | =   | <   | <   | -   | -   | <   | <   |
| WPR                                                                                                                                                                                                                                               | >            | >   | >   | >   | >   | =   | <   | -   | -   | <   | <   |
| WTR                                                                                                                                                                                                                                               | >            | >   | >   | >   | >   | >   | =   | -   | -   | -   | <   |
| DUN                                                                                                                                                                                                                                               | >            | -   | >   | -   | -   | -   | -   | =   | <   | <   | <   |
| DSP                                                                                                                                                                                                                                               | >            | >   | >   | >   | -   | -   | -   | >   | =   | <   | <   |
| DPR                                                                                                                                                                                                                                               | >            | >   | >   | >   | >   | >   | -   | >   | >   | =   | <   |
| DTR                                                                                                                                                                                                                                               | >            | >   | >   | >   | >   | >   | >   | >   | >   | >   | =   |
| <b>Key</b><br>< <i>mode1</i> is weaker than <i>mode2</i><br>> <i>mode1</i> is stronger than <i>mode2</i><br>= <i>mode1</i> and <i>mode2</i> are the same<br>-     there is no relation of relative strength between <i>mode1</i> and <i>mode2</i> |              |     |     |     |     |     |     |     |     |     |     |

Table 5 - Compatibility of lock modes

| <i>mode1</i>                                                                                                          | <i>mode2</i> |     |     |     |     |     |     |     |     |     |     |
|-----------------------------------------------------------------------------------------------------------------------|--------------|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
|                                                                                                                       | RUN          | RSP | WUN | WSP | RPR | WPR | WTR | DUN | DSP | DPR | DTR |
| RUN                                                                                                                   | yes          | yes | yes | yes | yes | yes | yes | yes | yes | yes | yes |
| RSP                                                                                                                   | yes          | yes | yes | yes | yes | yes | yes | no  | no  | no  | no  |
| WUN                                                                                                                   | yes          | yes | yes | yes | no  | no  | no  | yes | yes | no  | no  |
| WSP                                                                                                                   | yes          | yes | yes | yes | no  | no  | no  | no  | no  | no  | no  |
| RPR                                                                                                                   | yes          | yes | no  | no  | yes | no  | no  | no  | no  | no  | no  |
| WPR                                                                                                                   | yes          | yes | no  | no  | no  | no  | no  | no  | no  | no  | no  |
| WTR                                                                                                                   | yes          | yes | no  | no  | no  | no  | no  | no  | no  | no  | no  |
| DUN                                                                                                                   | yes          | no  | yes | no  | no  | no  | no  | no  | no  | no  | no  |
| DSP                                                                                                                   | yes          | no  | yes | no  | no  | no  | no  | no  | no  | no  | no  |
| DPR                                                                                                                   | yes          | no  | no  | no  | no  | no  | no  | no  | no  | no  | no  |
| DTR                                                                                                                   | yes          | no  | no  | no  | no  | no  | no  | no  | no  | no  | no  |
| <b>Key</b><br>yes <i>mode1</i> and <i>mode2</i> are compatible<br>no <i>mode1</i> and <i>mode2</i> are not compatible |              |     |     |     |     |     |     |     |     |     |     |

Table 6 - Implicit promotion of explicit lock of mode *mode1* to *mode2*

|                                                                                   | <i>mode2</i> |     |     |     |     |     |     |     |
|-----------------------------------------------------------------------------------|--------------|-----|-----|-----|-----|-----|-----|-----|
| <i>mode1</i>                                                                      | RUN          | WUN | RPR | WPR | WTR | DUN | DPR | DTR |
| <b>RUN</b>                                                                        | no           | WUN | no  | WUN | WUN | DUN | DUN | DTR |
| <b>RSP</b>                                                                        | no           | WSP | no  | WSP | WSP | DSP | DSP | DTR |
| <b>WUN</b>                                                                        | no           | no  | no  | no  | no  | DUN | DUN | DTR |
| <b>WSP</b>                                                                        | no           | no  | no  | no  | no  | DSP | DSP | DTR |
| <b>RPR</b>                                                                        | no           | WPR | no  | WPR | WPR | DPR | DPR | DTR |
| <b>WPR</b>                                                                        | no           | no  | no  | no  | no  | DPR | DPR | DTR |
| <b>WTR</b>                                                                        | -            | -   | -   | -   | no  | -   | -   | DTR |
| <b>DUN</b>                                                                        | no           | no  | no  | no  | no  | no  | no  | DTR |
| <b>DSP</b>                                                                        | no           | no  | no  | no  | no  | no  | no  | DTR |
| <b>DPR</b>                                                                        | no           | no  | no  | no  | no  | no  | no  | DTR |
| <b>DTR</b>                                                                        | -            | -   | -   | -   | no  | -   | -   | DTR |
| <b>Key</b><br>no        there is no promotion<br>-        the case does not apply |              |     |     |     |     |     |     |     |

Table 7 - Promotion of *mode1* to *mode2*: other cases

|              | <i>mode2</i> |     |     |     |     |     |     |     |     |     |     |
|--------------|--------------|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| <i>mode1</i> | RUN          | RSP | WUN | WSP | RPR | WPR | WTR | DUN | DSP | DPR | DTR |
| <b>RUN</b>   | no           | RSP | WUN | WSP | RPR | WPR | WTR | DUN | DSP | DPR | DTR |
| <b>RSP</b>   | no           | no  | WSP | WSP | RPR | WPR | WTR | DSP | DSP | DPR | DTR |
| <b>WUN</b>   | no           | WSP | no  | WSP | WPR | WPR | WTR | DUN | DSP | DPR | DTR |
| <b>WSP</b>   | no           | no  | no  | no  | WPR | WPR | WTR | DSP | DSP | DPR | DTR |
| <b>RPR</b>   | no           | no  | WPR | WPR | no  | WPR | WTR | DPR | DPR | DPR | DTR |
| <b>WPR</b>   | no           | no  | no  | no  | no  | no  | WTR | DPR | DPR | DPR | DTR |
| <b>WTR</b>   | no           | no  | no  | no  | no  | no  | no  | DTR | DTR | DTR | DTR |
| <b>DUN</b>   | no           | DSP | no  | DSP | DPR | DPR | DTR | no  | DSP | DPR | DTR |
| <b>DSP</b>   | no           | no  | no  | no  | DPR | DPR | DTR | no  | no  | DPR | DTR |
| <b>DPR</b>   | no           | no  | no  | no  | no  | no  | DTR | no  | no  | no  | DTR |
| <b>DTR</b>   | no           | no  | no  | no  | no  | no  | no  | no  | no  | no  | no  |

**Table 8 - Default External Lock Modes for Object Resources**

| Activity class     | Default lock mode |        |                 |                 |
|--------------------|-------------------|--------|-----------------|-----------------|
|                    | Read              | Write  |                 |                 |
|                    |                   | Update | Object creation | Object deletion |
| <b>UNPROTECTED</b> | RUN               | WUN    | WUN             | DUN             |
| <b>PROTECTED</b>   | RPR               | WPR    | WPR             | DPR             |
| <b>TRANSACTION</b> | RPR               | WTR    | DTR             | DTR             |

**Table 9 - Default External Lock Modes for Link Resources**

| Activity class     | Default lock mode |        |               |               |
|--------------------|-------------------|--------|---------------|---------------|
|                    | Read              | Write  |               |               |
|                    |                   | Update | Link creation | Link deletion |
| <b>UNPROTECTED</b> | RUN               | WUN    | WUN           | WUN           |
| <b>PROTECTED</b>   | RPR               | WPR    | WPR           | WPR           |
| <b>TRANSACTION</b> | RPR               | WTR    | WTR           | WTR           |

## 16.2 Concurrency and integrity control operations

### 16.2.1 ACTIVITY\_ABORT

ACTIVITY\_ABORT (  
)

ACTIVITY\_ABORT terminates the current activity of the calling process and discards uncommitted updates. The following actions are performed in order:

- The activity status and activity start termination time of the current activity of the calling process are set to ABORTING and the system current time respectively.

This implies abnormal termination of any process P which was initiated in the context of the current activity and the execution of which has been started but not yet terminated (i.e. a process in one of the states RUNNING, SUSPENDED, and STOPPED), in the same way as by calling PROCESS\_TERMINATE (P, ACTIVITY\_ABORTED).

ACTIVITY\_ABORT waits until all those processes have terminated (i.e. have the state TERMINATED). If while in this phase ACTIVITY\_ABORT is interrupted, either because the time-out period for the calling process has expired, or because another process has called PROCESS\_INTERRUPT\_OPERATION for the calling process, then the current activity and

the associated resources are not affected. In particular, the activity status and activity start termination time of the activity are reset to their previous values.

- If the current activity is a transaction, all updates made on behalf of the activity to the resources acquired with WRITE\_TRANSACTIONED or DELETE\_TRANSACTIONED external mode lock since the establishing of the locks or the promotion of their external mode to WRITE\_TRANSACTIONED (or DELETE\_TRANSACTIONED for the locks which were directly established or promoted to this mode) are discarded.
- All the locks held by the activity are discarded. This results in the deletion of the "lock" and "locked\_by" links associated with those locks.
- The activity status and activity start termination time of the activity are set to ABORTED and the current system time respectively. The activity remains in existence until the calling process is deleted.

As a result of these actions, the activity on whose behalf the aborted activity was initiated becomes the calling process's current activity.

### Errors

ACTIVITY\_WAS\_NOT\_STARTED\_BY\_CALLING\_PROCESS

ACTIVITY\_IS\_OPERATING\_ON\_A\_RESOURCE

## 16.2.2 ACTIVITY\_END

```
ACTIVITY_END (
)
```

ACTIVITY\_END terminates the current activity of the calling process normally. The effect of this operation is immediately to commit all outstanding updates in the context of the enclosing activities and to release all locks still held by the activity. The following actions are performed in order:

- The activity status and activity start termination time of the current activity of the calling process are set to COMMITTING and the current system time respectively.

The operation then waits until all the processes which were initiated on behalf of the activity and the execution of which has been started but not yet terminated (i.e. processes which are running, suspended, or stopped) have terminated.

If, while in this phase, the operation is interrupted, either because the time-out period defined for the calling process has expired or because another process has called PROCESS\_INTERRUPT\_OPERATION for the calling process, then the current activity and the associated resources are not affected. In particular, the activity status and activity start termination time of the activity are reset to their previous values.

- The locks still held by the activity are released. For locks established with external mode WRITE\_TRANSACTIONED or DELETE\_TRANSACTIONED, all the updates made to the locked resource on behalf of the activity since the establishment of the lock are committed in the context of the enclosing transactions. This means the WRITE\_TRANSACTIONED and DELETE\_TRANSACTIONED locks are inherited by the closest enclosing transaction if any, otherwise the modification are committed (i.e. modification can no longer be discarded)



and those locks are discarded (see 16.1.4 and 16.1.8). In any case, this results in the deletion of the "lock" and "locked\_by" links associated with those locks.

- The activity status and activity start termination time of the activity are set to COMMITTED and the current system time respectively.

The activity object remains in existence until the calling process is deleted. The activity on whose behalf the terminated activity was initiated becomes the calling process's current activity.

### Errors

ACTIVITY\_WAS\_NOT\_STARTED\_BY\_CALLING\_PROCESS

ACTIVITY\_IS\_OPERATING\_ON\_A\_RESOURCE

TRANSACTION\_CANNOT\_BE\_COMMITTED

## 16.2.3 ACTIVITY\_START

```
ACTIVITY_START (
    activity_class : Activity_class
)
```

ACTIVITY\_START creates a new activity of activity class *activity\_class*, nested within the current activity of the calling process.

The activity is created on the same volume as the calling process, with a "started\_activity" link to it from the calling process. The activity has the same mandatory labels and the same atomic and composite ACLs as the calling process.

A "nested\_in" link and a "nested\_activity" link are created between the new activity and the current activity of the calling process.

The activity class and activity start time of the new activity are set to *activity\_class* and the current system time respectively.

The new activity then becomes the current activity for the calling process.

### Errors

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_ACTIVITIES)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_ACTIVITIES\_PER\_PROCESS)

If the calling process has OWNER granted or denied:

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION (new activity)

VOLUME\_IS\_FULL (volume on which the calling process resides)

NOTE - The class of an activity influences system behaviour with respect to lock durations and the default external mode of implicit locks.

## 16.2.4 LOCK\_RESET\_INTERNAL\_MODE

```
LOCK_RESET_INTERNAL_MODE (
    object : Object_designator
)
```

LOCK\_RESET\_INTERNAL\_MODE resets to READ\_UNPROTECTED the internal mode of the lock associated with the object resource *object*.

As result, the internal lock mode of the associated "lock" link is set to the value READ\_UNPROTECTED.

### Errors

DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED (*object*, ATOMIC)

LOCK\_IS\_NOT\_EXPLICIT (*object*)

OBJECT\_IS\_NOT\_LOCKED (*object*)

OBJECT\_IS\_OPERATED\_ON (*object*, ATOMIC)

## 16.2.5 LOCK\_SET\_INTERNAL\_MODE

```
LOCK_SET_INTERNAL_MODE (
    object      : Object_designator,
    lock_mode   : Lock_internal_mode,
    wait_flag   : Boolean
)
```

LOCK\_SET\_INTERNAL\_MODE promotes the internal mode of the lock on the object resource designated by *object*.

If the required lock internal mode is weaker than the existing one, no action is performed.

If *lock\_mode* is not stronger than the internal mode of the lock currently held by the activity on *object*, then, whenever possible, the operation results in an explicit promotion of the internal mode of that lock to the weakest mode which is stronger than both the current internal mode and *lock\_mode*. E.g. if the current internal mode is READ\_PROTECTED and *lock\_mode* is WRITE\_UNPROTECTED then, if the operation succeeds, it results in the promotion of the internal mode of the existing lock to WRITE\_PROTECTED.

Let *new\_lock\_mode* be the actual value of this lock internal mode: either the specified value *lock\_mode* or the value derived from it by the above promotion rule; then LOCK\_SET\_INTERNAL\_MODE sets the internal lock mode of the associated "lock" link to *new\_lock\_mode*.

In case of conflict between the required internal mode and other concurrent acquisitions of the resources in the concerned domain of the object resource *object* (see 16.1.7), the behaviour of the operation depends on the value of *wait\_flag*:

- **true** : the operation waits on the resource until it acquires the resource or until the operation is interrupted or until the process is terminated, whichever comes first;
- **false**: the operation does not wait on the resource and the operation fails with the error condition LOCK\_INTERNAL\_MODE\_CANNOT\_BE\_CHANGED and has no effect.

### Errors

DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED (*object*, ATOMIC)

LOCK\_INTERNAL\_MODE\_CANNOT\_BE\_CHANGED (*object*, *lock\_mode*)

LOCK\_IS\_NOT\_EXPLICIT (*object*)

LOCK\_MODE\_IS\_TOO\_STRONG (*lock\_mode*, *object*)

OBJECT\_IS\_NOT\_LOCKED (*object*)

## 16.2.6 LOCK\_SET\_OBJECT

```
LOCK_SET_OBJECT (
    object      : Object_designator,
    lock_mode   : Lock_set_mode,
    wait_flag   : Boolean,
    scope       : Object_scope
)
```

LOCK\_SET\_OBJECT either establishes a new lock on the object resource *object*, if *object* is not yet assigned to the current activity, or promotes an existing lock on *object* otherwise. If *scope* is COMPOSITE, LOCK\_SET\_OBJECT also does the same for the object resource of each component of the object *object*.

If *lock\_mode* is READ\_DEFAULT, WRITE\_DEFAULT or DELETE\_DEFAULT, the external mode of the lock is chosen according to the class of the current activity as shown in table 10.

**Table 10 - Interpretation of default lock modes**

| Activity class | READ_DEFAULT | WRITE_DEFAULT | DELETE_DEFAULT |
|----------------|--------------|---------------|----------------|
| UNPROTECTED    | RUN          | WUN           | DUN            |
| PROTECTED      | RPR          | WPR           | DPR            |
| TRANSACTION    | RPR          | WTR           | DTR            |

The locks are established or promoted as follows, where *resource* is the object resource *object*, and each of the object resources of the components of *object* if *scope* is COMPOSITE.

If the current activity of the calling process has a lock on *resource* then the lock's external mode is promoted to *lock\_mode*. If this is weaker than the lock's current external mode, the operation is successful but no action is performed.

The internal mode of a new lock is set to READ\_UNPROTECTED. If *lock\_mode* is not stronger than the lock's current external mode, then the operation results in an explicit promotion of the external mode of the lock according to the rule of explicit promotion of an external mode defined in 16.1.5.

Let *new\_lock\_mode* be the actual value of this lock mode: either the specified value *lock\_mode* or the value derived from it as defined by the above rule of explicit promoting an external mode.

If the current activity is enclosed in a transaction, LOCK\_SET\_OBJECT also results in an attempt to implicitly establish or to implicitly promote a lock on the object resource *resource* on behalf of the closest enclosing transaction. The external mode *implied\_lock\_mode* of this implied lock is derived from *lock\_mode* as defined in 16.1.6, and its internal mode *implied\_internal\_lock\_mode* is READ\_UNPROTECTED.

If new locks are to be established, for each of these locks, LOCK\_SET\_OBJECT creates a "lock" and a "locked\_by" link (each reversing the other) between the activity on behalf of which the lock is established (i.e. the current activity of the calling process or its closest enclosing transaction) and the locked object. The links remain in existence until the corresponding locks

are released or inherited. The keys of the "lock" and "locked\_by" links are implementation-dependent.

In this case LOCK\_SET\_OBJECT initializes the attributes of the new links as follows:

- Attributes of "lock" link of the current activity:
  - lock duration is set to LONG if the current activity is a transaction, otherwise it is set to SHORT;
  - lock explicitness is set to EXPLICIT;
  - lock internal mode and lock external mode are set to READ\_UNPROTECTED and to *new\_lock\_mode* respectively.
- Attributes of "lock" link of the enclosing transaction:
  - lock duration is set to LONG;
  - lock explicitness is set to IMPLICIT;
  - lock internal mode and lock external mode are set to *implied\_internal\_lock\_mode* and to *implied\_lock\_mode* respectively.

The locked resource of each of these links is not set and has its initial value which is the empty string.

If the locks are promoted then only the external lock mode and internal lock mode of the existing locks are modified; they are set as described above.

In case of conflict between the locks required on *object* or on any of its components (if *scope* is COMPOSITE) and other concurrent acquisitions of the resources in the corresponding concerned domain, the behaviour of the operation depends on the value of *wait\_flag*:

- **true:** no lock is established or promoted by the operation, and the operation waits on the resource until the resource becomes available, until the operation is interrupted, or until the process is terminated, whichever comes first;
- **false:** no lock is established or promoted by the operation, the operation does not wait on the resource, and the operation fails with the LOCK\_COULD\_NOT\_BE\_ESTABLISHED error condition.

If *scope* is COMPOSITE, then none of the locks that the operation is trying to establish or promote are established or promoted until all of them can be established or promoted.

### Errors

DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED (*object*, *scope*)

LOCK\_COULD\_NOT\_BE\_ESTABLISHED (*object*, *scope*)

LOCK\_MODE\_IS\_NOT\_ALLOWED (*lock\_mode*)

If *scope* is ATOMIC:

OBJECT\_IS\_ARCHIVED (*object*)

If *scope* is COMPOSITE:

OBJECT\_IS\_ARCHIVED (*object* or a component of *object*)

OBJECT\_IS\_INACCESSIBLE (*object*, *scope*)

## 16.2.7 LOCK\_UNSET\_OBJECT

```
LOCK_UNSET_OBJECT (
    object : Object_designator,
    scope  : Object_scope
)
```

LOCK\_UNSET\_OBJECT releases the lock established by the current activity of the calling process on the object resource *object*. If *scope* is COMPOSITE, LOCK\_UNSET\_OBJECT also does the same thing for the object resource of each component of *object*.

This results in the deletion of the "lock" and "locked\_by" links associated with the released lock or locks.

### Errors

DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED (*object*, *scope*)

LOCK\_IS\_NOT\_EXPLICIT (*object*)

If *scope* is ATOMIC:

OBJECT\_IS\_ARCHIVED (*object*)

If *scope* is COMPOSITE:

OBJECT\_IS\_ARCHIVED (*object* or a component of *object*)

OBJECT\_IS\_INACCESSIBLE (*object*, *scope*)

OBJECT\_IS\_OPERATED\_ON (*object*, *scope*)

UNLOCKING\_IN\_TRANSACTION\_IS\_FORBIDDEN

## 17 Replication

### 17.1 Replication concepts

#### 17.1.1 Replica sets

Replica\_set\_identifier = Natural

**sds** system:

replica\_set\_identifier: **natural**;

replica\_set\_directory: **child type of object with link**

known\_replica\_set: (**navigate**) **non\_duplicated existence link** (replica\_set\_identifier) **to** replica\_set **reverse** known\_replica\_set\_of;

replica\_sets\_of: **implicit link to common\_root reverse** replica\_sets;

**end** replica\_set\_directory;

replica\_set: **child type of object with link**

master\_volume: (**navigate**) **reference link to** administration\_volume **reverse** master\_volume\_of;

copy\_volume: (**navigate**) **reference link** (volume\_identifier) **to** administration\_volume **reverse** copy\_volume\_of;

known\_replica\_set\_of: **implicit link to** replica\_set\_directory **reverse** known\_replica\_set;

**end** replica\_set;

```

extend object type administration_volume with
link
    master_volume_of: (navigate) reference link (replica_set_identifier) to replica_set
    reverse master_volume;
    copy_volume_of: (navigate) reference link (replica_set_identifier) to replica_set reverse
    copy_volume;
end administration_volume;
end system;

```

The replica set directory represents the set of known replica sets. Each replica set has a unique replica set identifier which is assigned to the replica set on creation and uniquely identifies the replica set within the PCTE installation.

The replica set directory is the destination of a "replica\_sets" link from the common root. Each known replica set is the destination of a "known\_replica\_set" link from the replica set directory whose key is the replica set identifier of that replica set.

A replica set has exactly one master volume which is chosen when the replica set is created. It also has a set of copy volumes. Master and copy volumes are administration volumes. The key of a "copy\_volume" link is the volume identifier of its destination volume. The keys of "master\_volume" and "copy\_volume" links are the replica set identifiers of their destination replica sets.

A copy volume of a replicated set cannot also be the master volume for that same set.

### 17.1.2 Replicated objects

sds system:

```

extend object type replica_set with
link
    includes_object: (navigate) reference link (exact_identifier) to object reverse
    replicated_as_part_of;
end replica_set;

extend object type administration_volume with
link
    replica: (navigate) reference link (exact_identifier) to object reverse replica_on;
end administration_volume;

extend object type object with
link
    replicated_as_part_of: (navigate) implicit link to replica_set reverse includes_object;
    replica_on: implicit link to administration_volume reverse replica;
end object;
end system;

```

Objects are classified as *normal*, *master* or *copy*, according to the value of the replicated state (see 9.1.1):

- A *master* object has replicated state MASTER; it belongs to exactly one replica set, and it resides on the master volume of that replica set.
- A *copy* object has replicated state COPY; it belongs to exactly one replica set, and is a replica on a copy volume of that replica set, but does not reside on any volume (i.e. there is no "object\_on\_volume" link to it; see 11.1.1).

- A *normal* object has replicated state NORMAL; it belongs to no replica set, and can reside on any volume.

For each master object there may be a corresponding copy object (with the same exact identifier) on each of its replica set's copy volumes; for each copy object there is a corresponding master object on its replica set's master volume. Such a set of corresponding master and copy objects is called a *replicated object*.

Operations which modify a copy object are forbidden, and master objects can be modified only by processes for which PCTE\_REPLICATION is an effective security group.

The destinations of the "includes\_object" links leaving a replica set are called the objects *replicated as part of* that replica set. The key of an "includes\_object" link is the exact identifier of its destination object.

If a link is created to a replicated object then the link is created to the master object.

A replica set is replicated as part of itself.

The following objects cannot be replicated: processes, activities, pipes, devices, execution sites, volumes, message queues, audit files, and accounting logs.

#### NOTES

- 1 There is intended to be a copy of each object of a replicated set on each of the copy volumes of the replicated set.
- 2 The master and all copies of a replicated object are intended to be kept identical, except for the volume identifier, last access time, replicated state, composite last access time, composite last change time, composite last modification time and "replica\_on" and usage designation links. It is expected that system tools automatically propagate modifications and enforce convergence among the various copies in a PCTE installation. Instantaneous updating of all copies of a replicated object as the master evolves is not expected, so that the replication mechanism has to manage temporary inconsistencies among the various copies of the replicated objects, supporting suitable procedures for the propagation of updates.

### 17.1.3 Selection of an appropriate replica

sds system:

**extend object type** workstation with  
**link**

    replica\_set\_chosen\_volume: (**navigate**) **designation link** (replica\_set\_identifier) to  
    administration\_volume;

**end** workstation;

**end** system;

At each moment for each workstation W and replica set S, there is a unique volume V called the *chosen volume* of W for accessing S. This is defined as follows:

- it is an *explicitly chosen volume*. This explicit choice is modelled as a "replica\_set\_chosen\_volume" link from W to V with the replica set identifier of S as its key.
- it is an *implicitly chosen volume*. This implicit choice is made when there is no "replica\_set\_chosen\_volume" link from W with the replica set identifier of S as its key.

If a link destination resides on an administration volume this is considered to be a potential link to a replicated object.



If a link identifies an object *O* that is replicated as part of a replica set *S* then replication redirection may occur. *Replication redirection* means that the object reached or navigated through as the destination of a link (including "replica" links) is the copy object of *O* on the chosen volume for accessing *S* of the execution site of the calling process (see 17.1.3) if such a copy object exists, and the master object of *O* otherwise. There is an exception for service designation links in that replication redirection applies to the state of the object base at the time the link was created rather than when it is navigated through.

When an object is referenced using an internal object reference or a contents handle the designated object is always the one that was reached at the time the corresponding pathname evaluation was performed, regardless of whether a new local copy has been created or deleted since that evaluation.

An object replicated as part of a replica set *S* can only be updated by processes having PCTE\_REPLICATION privilege and running on a workstation whose chosen volume for accessing *S* is the master volume of *S*. If a link having referential integrity is created to a replicated object then the calling process must have PCTE\_REPLICATION privilege.

#### NOTES

- 1 Since "replica\_set\_chosen\_volume links" are designation links they may be replaced when the volumes that they designate become unavailable, thus allowing an alternative volume containing that replica set to be chosen.
- 2 A PCTE implementation may automatically decide which of the master or copy volumes of a replica set should be implicitly chosen for each workstation. This can be used to allow the PCTE implementation to minimize network load and to recover from machine or network failure.
- 3 A navigation to or from a replicated object should not fail because the master object is not accessible as long as the chosen volume of the workstation of the process on whose behalf the navigation is being performed contains a copy of that object.
- 4 Access to a direct component of an object is made as if the corresponding composition link is navigated.

### 17.1.4 Administration replica set

There is exactly one *administration replica set*. It has replica set identifier 0. Its master volume is the master administration volume (see 11.1.2). Each administration volume other than the master administration volume is a copy volume of the administration replica set.

The administration replica set includes the *predefined replicated objects*, representing certain system entities; each predefined replicated object has a master on the master administration volume and a copy on each other administration volume; a predefined replicated object may not have its replicated state changed. The predefined replicated objects are:

- the common root;
- the administrative objects.

NOTE - Copies of predefined replicated objects cannot be deleted.

## 17.2 Replication operations

### 17.2.1 REPLICA\_SET\_ADD\_COPY\_VOLUME

```
REPLICA_SET_ADD_COPY_VOLUME (
    replica_set      : Replica_set_designator,
    copy_volume      : Administration_volume_designator
)
```

REPLICA\_SET\_ADD\_COPY\_VOLUME adds *copy\_volume* to the set of copy volumes for the replica set *replica\_set*.

A "copy\_volume" link with key equal to the volume identifier of *copy\_volume* is created from *replica\_set* to *copy\_volume*. The key of its reverse link is the replica set identifier of *replica\_set*.

A copy of *replica\_set* is created in *copy\_volume*. A "replica" link with key equal to the exact identifier of *replica\_set* is created from *copy\_volume* to this copy object.

A read lock of the default mode is set of the master of *replica\_set*. Write locks of the default modes are set on the created "copy\_volume", "copy\_volume\_of", "replica" and "replica\_on" links and the created copy of *replica\_set*.

#### Errors

ACCESS\_ERRORS (*copy\_volume*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (master of *replica\_set*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (master of *replica\_set*, ATOMIC, READ, READ\_ATTRIBUTES)

ACCESS\_ERRORS (master of *replica\_set*, ATOMIC, READ, READ\_LINKS)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_REPLICATION)

REPLICA\_SET\_IS\_NOT\_KNOWN (*replica\_set*)

VOLUME\_IS\_ALREADY\_COPY\_VOLUME\_OF\_REPLICA\_SET (*replica\_set*, *copy\_volume*)

VOLUME\_IS\_MASTER\_VOLUME\_OF\_REPLICA\_SET (*replica\_set*, *copy\_volume*)

### 17.2.2 REPLICA\_SET\_CREATE

```
REPLICA_SET_CREATE (
    master_volume    : Administration_volume_designator,
    identifier        : Natural
)
    replica_set      : Replica_set_designator
```

REPLICA\_SET\_CREATE creates replica set *replica\_set* with master volume *master\_volume*. The newly created object resides on *master\_volume*, and is replicated as part of itself.

A "known\_replica\_set" link is created from the replica set directory to *replica\_set* with key *identifier*, together with its reverse link. *identifier* becomes the replica set identifier of the replica set.

A "master\_volume" link is created from *replica\_set* to *master\_volume*. The key of its reverse link is assigned the value *identifier*.

An "includes\_object" link is created from *replica\_set* to *replica\_set* with key equal to the exact identifier of *replica\_set*. Its reverse link is also created.

Write locks of the default modes are obtained on *replica\_set* and the created "known\_replica\_set", "known\_replica\_set\_of", "master\_volume", "master\_volume\_of", "includes\_object" and "replicated\_as\_part\_of" links.

#### Errors

ACCESS\_ERRORS (replica set directory, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*master\_volume*, ATOMIC, MODIFY, APPEND\_LINKS)

LINK\_EXISTS (replica set directory, "known\_replica\_set" link from replica set directory with key *identifier*.)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_REPLICATION)

### 17.2.3 REPLICA\_SET\_REMOVE

```
REPLICA_SET_REMOVE (
    replica_set : Replica_set_designator
)
```

REPLICA\_SET\_REMOVE removes replica set *replica\_set* from the replicated set directory.

The "master\_volume" link from *replica\_set* to its master volume is deleted, together with its reverse link.

The "includes\_object" link from *replica\_set* to *replica\_set* is deleted, together with its reverse link.

The "known\_replica\_set" link from the replica set directory to *replica\_set* is deleted, together with its reverse link. If this link is the last existence link leading to *replica\_set*, *replica\_set* is deleted.

Write locks of the default modes are set on *replica\_set* and the deleted "known\_replica\_set", "known\_replica\_set\_of", "master\_volume", "master\_volume\_of", "includes\_object" and "replicated\_as\_part\_of" links.

#### Errors

ACCESS\_ERRORS (replica set directory, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (*replica\_set*, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (*replica\_set*, ATOMIC, CHANGE, WRITE\_IMPLICIT)

ACCESS\_ERRORS (master volume of *replica\_set*, ATOMIC, MODIFY, WRITE\_LINKS)

OBJECT\_HAS\_LINKS\_PREVENTING\_DELETION (*replica\_set*)

OBJECT\_IS\_IN\_USE\_FOR\_DELETE (*replica\_set*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_REPLICATION)

REPLICA\_SET\_HAS\_COPY\_VOLUMES (*replica\_set*)

REPLICA\_SET\_IS\_NOT\_EMPTY (*replica\_set*)

REPLICA\_SET\_IS\_NOT\_KNOWN (*replica\_set*)

### 17.2.4 REPLICA\_SET\_REMOVE\_COPY\_VOLUME

```

REPLICA_SET_REMOVE_COPY_VOLUME (
    replica_set      : Replica_set_designator,
    copy_volume      : Administration_volume_designator
)

```

REPLICA\_SET\_REMOVE\_COPY\_VOLUME removes *copy\_volume* from the set of copy volumes of the replica set *replica\_set*. The copy of *replica\_set* on *copy\_volume* is deleted.

The "copy\_volume" link with key equal to the volume identifier of *copy\_volume* and its reverse link are deleted.

The link of type "replica" from *copy\_volume* to the copy of *replica\_set* on *copy\_volume* is deleted, as is its reverse link.

Write locks of the default modes are set on the deleted "copy\_volume", "copy\_volume\_of", "replica" and "replica\_on" links and the deleted copy of *replica\_set*.

#### Errors

ACCESS\_ERRORS (master of *replica\_set*, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (deleted copy of *replica\_set*, ATOMIC, MODIFY)

ACCESS\_ERRORS (*copy\_volume*, ATOMIC, MODIFY, WRITE\_LINKS)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_REPLICATION)

REPLICA\_SET\_COPY\_IS\_NOT\_EMPTY (*replica\_set*, *copy\_volume*)

REPLICA\_SET\_IS\_NOT\_UNKNOWN (*replica\_set*)

VOLUME\_IS\_NOT\_COPY\_VOLUME\_OF\_REPLICA\_SET (*replica\_set*, *copy\_volume*)

### 17.2.5 REPLICATED\_OBJECT\_CREATE

```

REPLICATED_OBJECT_CREATE (
    replica_set      : Replica_set_designator,
    object           : Object_designator
)

```

REPLICATED\_OBJECT\_CREATE converts the normal object *object* to a master object belonging to replica set *replica\_set*. The replicated state of *object* is set to MASTER.

A "replica" link is created from the master volume of *replica\_set* to *object* with key equal to the exact identifier of *object*. Its reverse link is also created to the master volume.

An "includes\_object" link is created from *replica\_set* to *object* with key equal to the exact identifier of *object*. Its reverse link is also created.

A write lock of the default mode is obtained on *object* and a write lock of the default mode is obtained on the created "replica" link.

#### Errors

ACCESS\_ERRORS (*object*, ATOMIC, CHANGE, CONTROL\_OBJECT)

ACCESS\_ERRORS (*object*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

ACCESS\_ERRORS (*replica\_set*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (master volume of *replica\_set*, ATOMIC, MODIFY, APPEND\_LINKS)

OBJECT\_IS\_NOT\_ON\_MASTER\_VOLUME\_OF\_REPLICA\_SET (*replica\_set*, *object*)

OBJECT\_IS\_REPLICATED (*object*)

OBJECT\_IS\_NOT\_REPLICABLE (*object*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_REPLICATION)

REPLICA\_SET\_IS\_NOT\_KNOWN (*replica\_set*)

The following implementation-dependent errors may be raised for any object *X* with a link to *object*:

VOLUME\_IS\_NOT\_MOUNTED (*X*, ATOMIC)

VOLUME\_IS\_READ\_ONLY (*X*, ATOMIC)

### 17.2.6 REPLICATED\_OBJECT\_DELETE\_REPLICA

```

REPLICATED_OBJECT_DELETE_REPLICA (
    object          : Object_designator,
    copy_volume     : Administration_volume_designator
)

```

REPLICATED\_OBJECT\_DELETE\_REPLICA deletes the copy object *object* from the volume *copy\_volume*. The "replica" link leading to *object* and its reverse "replica\_on" link are deleted.

If the copy object has contents and this is currently opened by one or more processes, the deletion of the contents is postponed until all processes have closed the contents; i.e. the object is no longer accessible for example using internal object references or for replication redirection, but an operation using a contents handle to access its contents is not affected by the deletion until the contents handle is closed.

Write locks of the default mode are obtained on *object* and on the deleted replica link.

#### Errors

ACCESS\_ERRORS (*copy\_volume*, ATOMIC, MODIFY, WRITE\_LINKS)

OBJECT\_IS\_NOT\_REPLICATED\_ON\_VOLUME (*object*, *copy\_volume*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_REPLICATION)

OBJECT\_IS\_PREDEFINED\_REPLICATED (*object*)

OBJECT\_IS\_A\_REPLICA\_SET (*object*)

STATIC\_CONTEXT\_IS\_IN\_USE (*object*)

The following implementation-dependent error may be raised:

ACCESS\_ERRORS (master of *object*, ATOMIC, CHANGE)

### 17.2.7 REPLICATED\_OBJECT\_DUPLICATE

```

REPLICATED_OBJECT_DUPLICATE (
    object          : Object_designator,
    volume          : Administration_volume_designator,
    copy_volume     : Administration_volume_designator
)

```

If *volume* and *copy\_volume* are the same, then REPLICATED\_OBJECT\_DUPLICATE has no effect.

If *volume* and *copy\_volume* are not the same, and a copy of *object* does not already exist in *copy\_volume*, a copy is created and a "replica" link, keyed by the exact identifier of *object*, is created from *copy\_volume* to the new copy, together with its reverse "replica\_on" link.

If *volume* and *copy\_volume* are not the same, and a copy of *object* already exists in *copy\_volume*, the copy in *copy\_volume* is updated as defined below.

On completion, the atomic object of the copy in *copy\_volume* is identical to the atomic object of *object* except for the following:

- The volume identifier of the copy object is set to the volume identifier of *copy\_volume*.
- The last access time of the copy object is set to the value of the system clock at the time of call.
- The destination of its "replica\_on" link is *copy\_volume*.
- The replicated state of the copy object is set to COPY.
- Usage designation links are not copied to the copy in *copy\_volume*.
- If *object* has contents and there is a copy of *object* in *copy\_volume*, the effect is that of CONTENTS\_TRUNCATE followed by CONTENTS\_WRITE with the contents of *object*.

A write lock of the default mode is obtained on the copy object, and a read lock of the default mode is obtained in *object*. Write locks of the default mode are obtained on the "replica" link and its reverse "replica\_on" link, if created.

## Errors

ACCESS\_ERRORS (copy of *object* on *volume*, ATOMIC, READ)

If a new "replica" link is created:

ACCESS\_ERRORS (*copy\_volume*, ATOMIC, MODIFY, APPEND\_LINKS)

OBJECT\_IS\_NOT\_REPLICATED\_ON\_VOLUME (*object*, *volume*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_REPLICATION)

REPLICATED\_COPY\_IS\_IN\_USE (*object*)

STATIC\_CONTEXT\_IS\_IN\_USE (*object*)

VOLUME\_IS\_NOT\_MASTER\_OR\_COPY\_VOLUME\_OF\_REPLICA\_SET (*volume*, replica set of *object*)

VOLUME\_IS\_NOT\_COPY\_VOLUME\_OF\_REPLICA\_SET (*copy\_volume*, replica set of *object*)

VOLUME\_IS\_MASTER\_VOLUME\_OF\_REPLICA\_SET (*copy\_volume*, replica set of *object*)

The following implementation-dependent error may be raised:

ACCESS\_ERRORS (master of *object*, ATOMIC, CHANGE)

## NOTES

1 REPLICATED\_OBJECT\_DUPLICATE causes a copy of the atomic object of *object* to exist as the atomic object of the copy object in the volume *copy\_volume*. The copy object has the same components as *object*, but the components are not copied.

2 Updates to copy objects by this operation are subject to transaction rollback.

### 17.2.8 REPLICATED\_OBJECT\_REMOVE

```

REPLICATED_OBJECT_REMOVE (
    object : Object_designator
)

```

REPLICATED\_OBJECT\_REMOVE removes the master object *object* from the replica set to which it belongs by changing it into a normal object residing on the master volume of this replica set. The replicated state of *object* is set to NORMAL, and the "replica" and "includes\_object" links leading to the object are deleted, together with their reverse "replica\_on" and "replicated\_as\_part\_of" links.

Write locks of the default mode are obtained on the deleted "replica", "replica\_on", "includes\_object" and "replicated\_as\_part\_of" links, and on *object*.

#### Errors

ACCESS\_ERRORS (master volume of *object*, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (*object*, ATOMIC, CHANGE, CONTROL\_OBJECT)

ACCESS\_ERRORS (*object*, ATOMIC, CHANGE, WRITE\_IMPLICIT)

ACCESS\_ERRORS (replica set containing *object*, ATOMIC, MODIFY, WRITE\_LINKS)

OBJECT\_HAS\_COPIES (*object*)

OBJECT\_IS\_A\_REPLICA\_SET (*object*)

OBJECT\_IS\_NOT\_MASTER\_REPLICATED\_OBJECT (*object*)

OBJECT\_IS\_PREDEFINED\_REPLICATED (*object*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_REPLICATION)

The following implementation-dependent errors may be raised for any object X with a link to *object*:

VOLUME\_IS\_NOT\_MOUNTED (X, ATOMIC)

VOLUME\_IS\_READ\_ONLY (X, ATOMIC)

### 17.2.9 WORKSTATION\_SELECT\_REPLICA\_SET\_VOLUME

```

WORKSTATION_SELECT_REPLICA_SET_VOLUME (
    station : Workstation_designator,
    replica_set : Replica_set_designator,
    volume : Administration_volume_designator
)

```

WORKSTATION\_SELECT\_REPLICA\_SET\_VOLUME selects *volume* as the chosen volume for accesses to *replica* set by *station*.

If *station* has a "replica\_set\_chosen\_volume" link whose key is the replica set identifier of *replica\_set*, that link is first deleted.

A "replica\_set\_chosen\_volume" link with a key equal to the replica set identifier of *replica\_set* is then created from *station* and leading to *volume*.

A write lock of the default mode is created on the "replica\_set\_chosen\_volume" link.



**Errors**

If *station* already has a chosen volume for accesses to *replica\_set*:

ACCESS\_ERRORS (*station*, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (*station*, ATOMIC, MODIFY, APPEND\_LINKS)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_REPLICATION)

REPLICA\_SET\_IS\_NOT\_KNOWN (*replica\_set*)

VOLUME\_IS\_NOT\_MASTER\_OR\_COPY\_VOLUME\_OF\_REPLICA\_SET (*replica\_set*,  
*volume*)

**17.2.10 WORKSTATION\_UNSELECT\_REPLICA\_SET\_VOLUME**

```
WORKSTATION_UNSELECT_REPLICA_SET_VOLUME (
    station      : Workstation_designator,
    replica_set  : Replica_set_designator
)
```

WORKSTATION\_UNSELECT\_REPLICA\_SET\_VOLUME deletes the "replica\_set\_chosen\_volume" link from *station* whose key is the replica set identifier of *replica\_set*.

A write lock of the default mode is created on the deleted "replica\_set\_chosen\_volume" link.

**Errors**

ACCESS\_ERRORS (*station*, ATOMIC, MODIFY, WRITE\_LINKS)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_REPLICATION)

REPLICA\_SET\_IS\_NOT\_KNOWN (*replica\_set*)

WORKSTATION\_HAS\_NO\_CHOICE\_OF\_VOLUME\_FOR\_REPLICA\_SET (*station*,  
*replica\_set*)

**18 Network connection****18.1 Network connection concepts****18.1.1 Execution sites**

**sds** system:

execution\_site\_directory: **child type of object with link**

known\_execution\_site: **non\_duplicated existence link** (execution\_site\_identifier) to  
execution\_site;

execution\_sites\_of: **implicit link to common\_root reverse** execution\_sites;

**end** execution\_site\_directory;

execution\_site: **child type of object with link**

running\_process: (**navigate**) **non\_duplicated designation link** (number) to process;

**end** execution\_site;

**end** system;

The execution site identifier is assigned to the execution site on creation and uniquely identifies the execution site within the PCTE installation during its existence.

The destinations of the "running\_process" links, if any, are the processes running on the workstation (see 13.1.4).

The execution site directory is an administrative object (see 9.1.2).

NOTE - An execution site is either a workstation (see 18.1.2) or a foreign system (see 18.1.3).

### 18.1.2 Workstations

Work\_status = **set of** Work\_status\_item

Work\_status\_item = ACTIVITY\_REMOTE\_LOCKS | ACTIVITY\_LOCAL\_LOCKS |  
TRANSACTION\_REMOTE\_LOCKS | TRANSACTION\_LOCAL\_LOCKS |  
QUEUE\_REMOTE | QUEUE\_LOCAL | RECEIVE\_REMOTE | RECEIVE\_LOCAL |  
CHILD\_REMOTE | CHILD\_LOCAL

Requested\_connection\_status = LOCAL | CLIENT | CONNECTED

Connection\_status = Requested\_connection\_status | AVAILABLE

Workstation\_status = Connection\_status \* Work\_status

New\_administration\_volume ::

|                        |                     |
|------------------------|---------------------|
| FOREIGN_DEVICE         | : String            |
| ADMINISTRATION_VOLUME  | : Volume_identifier |
| VOLUME_CHARACTERISTICS | : String            |
| DEVICE                 | : Device_identifier |
| DEVICE_CHARACTERISTICS | : String            |

**sds** system:

workstation: **child type of** execution\_site **with**  
**attribute**

connection\_status: **(read) non\_duplicated enumeration** (LOCAL, CLIENT,  
AVAILABLE, CONNECTED) := LOCAL;

PCTE\_implementation\_name: **(read) non\_duplicated string**;

PCTE\_implementation\_release: **(read) non\_duplicated string**;

PCTE\_implementation\_version: **(read) non\_duplicated string**;

node\_name: **(read) non\_duplicated string**;

machine\_name: **(read) non\_duplicated string**;

**link**

controlled\_device: **(navigate) non\_duplicated existence link** (device\_identifier:  
**natural**) **to** device **reverse** device\_of;

associated\_administration\_volume: **(navigate) non\_duplicated designation link to**  
administration\_volume;

initial\_process: **non\_duplicated existence link** (number) **to** process;

outermost\_activity: **(navigate) non\_duplicated existence link** (number) **to** activity;

**end** workstation;

**end** system;

The work status consists of a number of independent work status items as follows (where *local* means residing on a volume mounted on a device controlled by this workstation, and *remote* means residing on a volume mounted on a device controlled by some other workstation):

- ACTIVITY\_REMOTE\_LOCKS: at least one non-transaction activity started on the workstation holds locks on remote objects,

- **ACTIVITY\_LOCAL\_LOCKS**: at least one non-transaction activity started on another workstation has locks on local objects,
- **TRANSACTION\_REMOTE\_LOCKS**: at least one transaction started on the workstation holds locks on remote objects,
- **TRANSACTION\_LOCAL\_LOCKS**: at least one transaction started on another workstation has locks on local objects,
- **QUEUE\_REMOTE**: at least one process on the workstation has a remote reserved message queue,
- **QUEUE\_LOCAL**: at least one process on a remote workstation has a message queue reserved on the workstation,
- **RECEIVE\_REMOTE**: at least one process on the workstation is waiting for the reception of a message from a remote message queue,
- **RECEIVE\_LOCAL**: at least one process on another workstation is waiting for the reception of a message from a local message queue,
- **CHILD\_REMOTE**: at least one process on the workstation has one or more untermiated remote child processes,
- **CHILD\_LOCAL**: at least one process on another workstation has one or more untermiated local child processes.

New administration volumes are used in **WORKSTATION\_CREATE**. The meaning of **FOREIGN\_DEVICE** is implementation-defined. It is a string used to designate a new physical resource (i.e. not yet represented by a device object).

A workstation A is a *client* of a workstation B if at least one of the following is true:

- a process running on A has started a child process on B which is not yet terminated;
- a process running on A is accessing (i.e. reading from, writing to, or navigating through) an object residing on a volume mounted on a device controlled by B;
- there is a service designation link to an object residing on a volume mounted on a device controlled by B from a process running on A;
- a process running on A has reserved a message queue whose associated message queue object resides on a volume mounted on a device managed by B.

Conversely, a workstation A is a *server* of a workstation B if B is a client of A.

The connection status denotes the status of the workstation with respect to other workstations of the PCTE installation. The values have the following meanings (for the definitions of client and server see below).

- **LOCAL** The workstation cannot be a client or a server for another workstation. It does not respond to a call of **WORKSTATION\_CONNECT** from another workstation.
- **AVAILABLE** The same as **LOCAL** except that the workstation responds to a connection request from another workstation.
- **CLIENT** The workstation can be a client but not a server of another workstation. It does not respond to a connection request from another workstation.

- CONNECTED The workstation can be a client or a server of another workstation.

The implementation name is the name of the particular implementation of PCTE running on the workstation; it is implementation-defined.

The implementation release identifies of the release of the PCTE implementation running on the workstation; it is implementation-defined.

The implementation version identifies of the version of the PCTE implementation running on the workstation; it is implementation-defined.

The node name provides the mechanism to communicate to the network, e.g. the local area network address; it is implementation-defined.

The machine name is the name of the particular machine type of the workstation; it is implementation-defined.

The controlled devices are also called the devices *controlled* by the workstation. Each of the devices is identified by a device identifier which is unique within the set of devices controlled by the workstation.

For the administration volume, see 11.1.2. A workstation object resides on the administration volume of the workstation.

For the initial process of the workstation, see 13.1.5.

For the outermost activity of the workstation, see 16.1.1.

For the associated accounting log, see 22.1.2.

A workstation is *busy* if it has connection status CONNECTED and is a server of another workstation, or has connection status CLIENT and is a client of another workstation.

Within an operation, the *local* workstation is the workstation on which the calling process is executed.

#### NOTES

1 The normal situation in a PCTE installation is one of the following, though abnormal situations may occur:

- all workstations with connection status LOCAL;
- one workstation with connection status AVAILABLE, all other workstations with connection status LOCAL;
- two or more workstations with connection status CONNECTED or CLIENT, all other workstations with connection status LOCAL;
- all workstations with connection status CONNECTED or CLIENT.

2 In some implementations a workstation may have more than one "initial\_process" or "outermost\_activity" link. Only the destinations with the highest key are the initial process and outermost activity of the workstation, respectively. Destinations with other keys are remnants from previous sessions which allow an implementation-dependent tool to perform actions following workstation or system failure.

3 A workstation may access an administration volume shared with another workstation even if its connection status is LOCAL.

### 18.1.3 Foreign systems

**sds** system:

foreign\_system: **child type of** execution\_site **with**  
**attribute**

system\_class: **enumeration** (FOREIGN\_DEVICE, BARE\_MACHINE,  
HAS\_EXECUTIVE\_SYSTEM, SUPPORTS\_IPC\_AND\_CONTROL,  
SUPPORTS\_MONITOR) := BARE\_MACHINE;

**end** foreign\_system;

**end** system;

The system class indicates the level of interaction which is supported between PCTE processes and foreign processes started on the foreign system, as follows.

- FOREIGN\_DEVICE The foreign system can be used only as the foreign system for operations defined in 18.3.
- BARE\_MACHINE The foreign system is a bare machine executing no code other than the software under development. The only permitted operation by a PCTE process is PROCESS\_CREATE. Any further communication is prevented by the absence of any communication agent on the foreign system
- HAS\_EXECUTIVE\_SYSTEM The foreign system is a foreign executive system which accepts the creation, starting, and termination of processes on it, and can signal the end of their execution to the creating host process.
- SUPPORTS\_IPC\_AND\_CONTROL As for HAS\_EXECUTIVE\_SYSTEM and can also support at least the message queue mechanisms represented by the operations (see clause 14) MESSAGE\_RECEIVE\_NO\_WAIT, MESSAGE\_RECEIVE\_WAIT, MESSAGE\_SEND\_NO\_WAIT, and MESSAGE\_SEND\_WAIT, and the process control mechanisms such as process suspension and resumption.
- SUPPORTS\_MONITOR As for SUPPORTS\_IPC\_AND\_CONTROL and can also support the monitoring operations of 13.5.

NOTE - On a foreign system of system class, BARE\_MACHINE, PROCESS\_CREATE is intended to download the process but not to start its execution.

### 18.1.4 Network partitions

The execution sites of a PCTE installation may be connected together to share resources.

At any time, a set of workstations of a PCTE installation, each of which is running PCTE, and which are connected together, is called a *network partition*. The connection status of each workstation in the network partition controls the use which one workstation may make of resources controlled by another (see below).

An implementation may impose restrictions on the sets of workstations which can form a network partition. In particular, an implementation may or may not allow any single isolated workstation to be a network partition, and an implementation may or may not allow more than one network partition to exist at the same time (in this case the sets of workstations in the network partitions are disjoint since the connectedness relation is transitive).

The specification of the abstract operations must always be understood to be in the context of the calling process's network partition; for example, "an object is not accessible" must always be understood to mean that the object is not accessible within the calling process's network partition; the object may be accessible in some other network partition.

A network partition may rejoin other partitions by implementation-defined means so that workstations in the partition are now accessible to workstations in other partitions. Although the time at which the network failure is detected may be variable, the network failure must be detected in all partitions before the network partitions can be rejoined.

### 18.1.5 Accessibility

In order for an operation to operate on an entity, that entity must be *accessible*.

The rules for accessibility are as follows; except as given by these rules, no entity in a PCTE installation is accessible:

- the local workstation is accessible;
- if the local workstation has connection status CLIENT or CONNECTED, all workstations in the same network partition which have connection status CONNECTED are accessible;
- processes executing on accessible workstations are accessible;
- devices controlled by accessible workstations are accessible;
- volumes mounted on accessible devices are accessible;
- objects residing on accessible volumes, or which are replicas on accessible administration volumes, and their direct attributes, direct outgoing links, contents, and associated sequences of messages are accessible;
- the atomic object associated with an accessible object is accessible;
- locks on resources residing on accessible volumes are accessible;
- the accessibility of a foreign system is implementation-defined;
- a process on an accessible foreign system is accessible;
- the accessibility of a file residing on a foreign system is implementation-defined.

A PCTE implementation may arrange that an operation succeeds even though some entity which is used is not accessible according to the above rules. Otherwise an error is raised to indicate that the entity is not accessible, and the above rules indicate what must be done by the user to make the entity accessible.

An object is defined as *unreachable* if operations fail to access it because it is not accessible.

When a workstation A ceases to be a client of a workstation B the following actions occur:

- All child processes running on B started by unreachable processes running on A are terminated in the same way as by calling PROCESS\_TERMINATE (child process, FORCED\_TERMINATION).
- All objects residing on volumes mounted on devices controlled by B and with contents opened by unreachable processes running on A are closed.

- All locks on objects or links residing on volumes mounted on devices controlled by B, or on deleted objects which resided on such volumes, and held by unreachable activities started by processes running on A, are treated as for activity abortion. If such a lock had external mode WTR or DTR then any changes are rolled back.
- Operations being executed by processes running on A accessing unreachable objects residing on volumes mounted on devices controlled by B may fail with the error condition OPERATION\_IS\_INTERRUPTED.
- All message queues residing on volumes mounted on devices controlled by B and reserved for unreachable processes running on A are unreserved, handlers on those message queues are disabled, and notifiers on those message queues are deleted.
- Usage designation links from objects residing on volumes mounted on devices controlled by B to unreachable processes running on A are deleted.

When a station A *ceases to be a server* of a station B the following actions occur:

- All unreachable objects residing on volumes mounted on devices controlled by A and with contents opened by processes running on B are closed. Any subsequent reads and writes to the contents fail.
- All locks on unreachable objects or links residing on volumes mounted on devices controlled by A held by activities started by processes running on B are released. These activities are not aborted immediately in their own workstations. However, any attempt to end a transaction activity which held such locks with external mode RPR, WTR or DTR results in the error TRANSACTION\_CANNOT\_BE\_COMMITTED.
- Operations being executed by processes running on B accessing unreachable objects which are residing on volumes mounted on devices controlled by A may fail with the error condition OPERATION\_IS\_INTERRUPTED.
- All unreachable message queues residing on volumes mounted on devices controlled by A and reserved for processes running on other workstations are unreserved.
- All "process\_waiting\_for" links from processes running on B to unreachable objects residing on volumes mounted on devices controlled by A are deleted, and the corresponding waiting operations in the processes fail with the error condition OBJECT\_IS\_INACCESSIBLE.
- Service designation links from processes running on B to unreachable objects residing on volumes mounted on devices controlled by A are deleted.
- For each "notifier" link whose origin message queue resides on a volume mounted on a device controlled by B and whose destination resides on a volume mounted on a device controlled by A, a NOT\_ACCESSIBLE\_MSG message is sent to its origin message queue.

If a network failure is detected by a workstation A then an inaccessible workstation B normally ceases to be a client of A and also ceases to be a server of A, and A ceases to be a client or server of B. However if B fails to detect the network failure (e.g. because it was only transient and connection at the communications protocol level has been restored) then any subsequent attempt by B to act as a server or client of A results in A responding so as to cause B to cease to be a server or client of A.



### 18.1.6 Workstation closedown

An *orderly closedown* of a workstation occurs only when all the descendant processes of the initial process associated with the workstation and any other processes executing on the workstation, with the exception of the initial process, are terminated. The outermost activity is terminated.

If a workstation is improperly terminated or fails, this is termed *abnormal closedown*. On abnormal closedown of a workstation, the following actions are taken:

- Each process P executing on the workstation at the time of workstation failure is terminated as by `PROCESS_TERMINATE (P, SYSTEM_FAILURE)`. In particular, all activities started by the processes are aborted.
- The outermost activity of the workstation is terminated abnormally.
- Contents of an implementation-dependent set of pipes managed by the workstation are lost. The messages and the values of the "reader\_waiting", "writer\_waiting", "space\_used", "message\_count", "last\_send\_time", and "last\_receive\_time" attributes of an implementation-dependent set of message queues managed by the workstation are lost.
- All locks on resources, residing on volumes mounted on devices controlled by the failed workstation, and held by activities started by processes executing on other workstations, are released as if the activities were abnormally terminated, and any updates performed under WTR or DTR locks are rolled back.
- All objects residing on volumes mounted on devices controlled by the failed workstation with contents opened by processes running on other workstations are closed.
- All message queues residing on volumes mounted on devices controlled by the failed workstation and reserved for processes running on other workstations are unreserved, handlers on those message queues are disabled, and notifiers on those message queues are deleted.
- Usage designation links from objects residing on volumes mounted on devices controlled by the failed workstation to processes running on other workstations are deleted.
- Updates to objects residing on volumes mounted on devices controlled by the failed workstation which had not been made permanent at the actual time of workstation failure are lost.
- Updates to files not under WTR or DTR locks, and to audit files and accounting logs residing on volumes mounted on devices controlled by the failed workstation are lost to an implementation-defined degree.
- The workstation connection status is set to LOCAL.

The terminated outermost activity and the terminated initial process remain.

Whether a workstation closes down in an abnormal or orderly manner, the contents of pipes, messages in message queues, and audit criteria are lost.

When a workstation is restarted after abnormal or orderly closedown, a new outermost activity and a new initial process are created, and the previous, terminated, outermost activity and initial process are not reused.

## NOTES

- 1 The actions described above at workstation abnormal closedown are performed by the implementation at some point between the failure and restarting the workstation.
- 2 As a consequence of terminating all the processes on abnormal closedown, any active activities are aborted.
- 3 After abnormal closedown the previous initial process may have components, i.e. be a tree of process objects.
- 4 The previously running initial process and the previously active outermost activity must be deleted explicitly. This means that it is possible that there are several "initial\_process" and several "outermost\_activity" links emanating from a workstation. However, only one "initial\_process" link leads to a running initial process and only one "outermost\_activity" link leads to an active outermost activity.
- 5 Implementations aiming for high security may wish to take special measures to ensure that workstation failure does not result in any loss of data written to the audit file.

## 18.2 Network connection operations

### 18.2.1 WORKSTATION\_CONNECT

```
WORKSTATION_CONNECT (
    status : Requested_connection_status
)
```

WORKSTATION\_CONNECT has no effect if the connection status of the local workstation is already the requested status *status*. Otherwise it connects the local workstation to the PCTE installation, and sets its connection status as follows:

- If *status* is CONNECTED and there is no other workstation in the PCTE installation available for connection, the connection status is set to AVAILABLE.
- If *status* is CLIENT and there is no other workstation in the PCTE installation available for connection, the connection status is unchanged.
- If *status* is CONNECTED or CLIENT and there is another workstation in the PCTE installation available for connection, the connection status is set to *status*. The connection status of all available workstations are automatically changed to CONNECTED.

There may be installation-defined procedures to be carried out before and after calling this operation; such procedures are outside the scope of this part of ISO/IEC 13719.

#### Errors

ACCESS\_ERRORS (the local workstation, ATOMIC, MODIFY, WRITE\_ATTRIBUTES)  
 CONNECTION\_IS\_DENIED  
 LAN\_ERROR\_EXISTS  
 PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_CONFIGURATION)  
 STATUS\_IS\_BAD (*status*)

## 18.2.2 WORKSTATION\_CREATE

```

WORKSTATION_CREATE (
    execution_site_identifier : Natural,
    administration_volume    : Volume_designator | New_administration_volume,
    access_mask              : Atomic_access_rights,
    node_name                : Text,
    machine_name             : Text
)

```

WORKSTATION\_CREATE creates a new workstation in the PCTE installation, as follows.

If *administration\_volume* is a volume designator:

- a "workstation" object *new\_station* is created on *administration\_volume* as destination of a new "known\_execution\_site" link from the execution site directory keyed by *execution\_site\_identifier*;
- an "object\_on\_volume" link is created from *existing\_administration\_volume* to *new\_station*, keyed by the exact identifier of *new\_station*.

If *administration\_volume* is a new administration volume with foreign device *foreign\_device*, administration volume *new\_administration\_volume*, volume characteristics *volume\_characteristics*, device *new\_device*, and device characteristics *device\_characteristics*:

- *foreign\_device* is interpreted in an implementation-defined way to specify a device containing a physical volume that has been prepared in an implementation-defined way to become a new administration volume.
- an "administration\_volume" object is created on the specified physical volume, with volume characteristics *volume\_characteristics*, as destination of a new "known\_volume" link from the volume directory, keyed by *new\_administration\_volume*;
- a "workstation" object *new\_station* is created on the new volume, as destination of a new "known\_execution\_site" link from the execution site directory keyed by *execution\_site\_identifier*;
- a "device\_supporting\_volume" object is created on the new volume, with device characteristics *device\_characteristics*, as destination of a new "controlled\_device" link from *new\_station*, keyed by *new\_device*;
- "object\_on\_volume" links are created from the new administration volume to itself, to *new\_station*, and to the new "device\_supporting\_volume" object, keyed by the exact identifiers of their destinations;
- the labels of the created device and of the created volume are set to the mandatory context of the calling process;
- a "mounted\_volume" link is created from the new device to the new administration volume;
- a "copy\_volume" link with key *volume\_identifier* is created to the newly created administration volume object, and is reversed by a "copy\_volume\_of" link with key '0' leading to the administration replica set. Copies of the master objects of the administration replica set, the common root, the "system" schema, and the administrative objects are created in the newly created administration volume.

In both cases:

- *access\_mask* is used in conjunction with the default atomic ACL and default owner of the calling process to define the atomic ACL and the composite ACL which are to be associated with the created objects (see 19.1.4);
- the labels of the created workstation are set to the mandatory context of the calling process;
- an "associated\_administration\_volume" link is created from *new\_station* to existing\_administration\_volume or new\_administration\_volume;
- an initial process is created for the workstation;
- the attributes of the new workstation are set as follows:
  - . the connection status is set to LOCAL;
  - . the PCTE implementation name, PCTE implementation release, and PCTE implementation version are the same as for the local workstation;
  - . the node name and machine name are set from the parameters *node\_name* and *machine\_name*.

If the auditing module is supported, there is at least one audit file for the new workstation (see 21.1.1), but auditing is initially disabled.

Write locks are obtained on the execution site directory, the volume directory, the created workstation and (if they are created) the new administration volume and device.

### Errors

ACCESS\_ERRORS (volume directory, ATOMIC, CHANGE, APPEND\_LINKS)

ACCESS\_ERRORS (execution site directory, ATOMIC, CHANGE, APPEND\_LINKS)

If *administration\_volume* is a volume designator:

ACCESS\_ERRORS (*administration\_volume*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (existing device, ATOMIC, CHANGE, APPEND\_IMPLICIT)

ACCESS\_ERRORS (administration replica set, ATOMIC, MODIFY, APPEND\_LINKS)

VOLUME\_IS\_UNKNOWN (*administration\_volume*)

CONTROL\_WOULD\_NOT\_BE\_GRANTED (*new\_station*)

If *administration\_volume* is a new administration volume:

FOREIGN\_DEVICE\_IS\_INVALID (*foreign\_device*)

VOLUME\_EXISTS (*new\_administration\_volume*)

VOLUME\_IDENTIFIER\_IS\_INVALID (*new\_administration\_volume*)

OBJECT\_OWNER\_VALUE\_WOULD\_BE\_INCONSISTENT\_WITH\_ATOMIC\_ACL

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_CONFIGURATION)

PROCESS\_IS\_IN\_TRANSACTION

WORKSTATION\_EXISTS (*execution\_site\_identifier*)

WORKSTATION\_IDENTIFIER\_IS\_INVALID (*execution\_site\_identifier*)

### NOTES

1 The new physical administration volume may need to be initialized by a system tool before this operation is invoked.

- 2 For bootstrapping reasons, this operation cannot apply to the first workstation of a PCTE installation.
- 3 The new workstation is created but is not yet initialized. It is an implementation-defined procedure which is responsible for starting the initial process of the new created workstation.
- 4 The ability to provide an existing administration volume is intended to cater for discless workstations and other cases of shared administration volumes.

### 18.2.3 WORKSTATION\_DELETE

```
WORKSTATION_DELETE (
    station : Workstation_designator
)
```

WORKSTATION\_DELETE deletes a workstation from the PCTE installation.

If the administration volume of *station* is mounted on a device which is controlled by another workstation (which implies that they share the same administration volume), the effect of the workstation deletion is the same as:

OBJECT\_DELETE (execution site directory, *execution\_site\_link*)

where *execution\_site\_link* is the "known\_execution\_site" link from the execution site directory to *station*.

If the administration volume of *station* is mounted on a device which is controlled by *station*, the workstation deletion is only possible if and only if:

- no other workstation has the same administration volume, i.e. there is only one "associated\_administration\_volume" link to the administration volume;
- the only objects residing on or which are replicas on the administration volume are:
  - . *station*;
  - . copies of the administration replica set;
  - . the administration volume;
  - . the "device\_supporting\_volume" object which is the destination of a "controlled\_device" link from *station* and the origin of the "mounted\_volume" link to the administration volume;
  - . terminated processes and activities;
- there are no reference, composition, or existence links from an object residing on another volume to the objects residing on the administration volume, except the "known\_volume" link from the volume directory to the administration volume, the "known\_execution\_site" link from the execution site directory to *station*, and "audit" links from *station*.

The objects residing on and which are replicas on the administration volume are deleted, the space previously occupied by the volume is freed, the "copy\_volume" link from the administration replica set to the administration volume, and the "known\_volume" link to the administration volume and the "known\_execution\_site" link to *station* are deleted. The administration volume is unmounted.

Write locks are obtained on the deleted workstation, the deleted administration volume, the deleted device supporting the administration volume, the administration replica set, and the

deleted links. These locks do not prevent the dismounting and deletion of the administration volume.

#### Errors

ACCESS\_ERRORS (execution site directory, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (volume directory, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (*station*, ATOMIC, CHANGE, WRITE\_IMPLICIT)

ACCESS\_ERRORS (*station*, ATOMIC, MODIFY, WRITE\_CONTENTS)

If the conditions hold for deletion of the "workstation" object *station*:

ACCESS\_ERRORS (*station*, COMPOSITE, MODIFY, DELETE)

OBJECT\_IS\_IN\_USE\_FOR\_DELETE (*station*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_CONFIGURATION)

PROCESS\_IS\_IN\_TRANSACTION

VOLUME\_HAS\_OTHER\_LINKS (*administration\_volume*)

VOLUME\_HAS\_OTHER\_OBJECTS (*administration\_volume*)

VOLUME\_HAS\_OBJECTS\_IN\_USE (*administration volume of station*)

WORKSTATION\_IS\_CONNECTED (*station*)

WORKSTATION\_IS\_UNKNOWN (*station*)

NOTE - Additional implementation-defined restrictions may be defined for this operation.

### 18.2.4 WORKSTATION\_DISCONNECT

```
WORKSTATION_DISCONNECT (
)
```

WORKSTATION\_DISCONNECT changes the connection status of the local workstation to LOCAL, unless the connection status is already LOCAL, in which case it has no effect.

#### Errors

ACCESS\_ERRORS (local workstation, ATOMIC, WRITE)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_CONFIGURATION)

WORKSTATION\_IS\_BUSY (local workstation)

### 18.2.5 WORKSTATION\_GET\_STATUS

```
WORKSTATION_GET_STATUS (
    station : [ Workstation_designator ]
)
    status : Workstation_status
```

WORKSTATION\_GET\_STATUS returns the current connection status and work status of *station* in *status*. If *station* is not supplied, the local workstation is assumed.

#### Errors

ACCESS\_ERRORS (*station*, ATOMIC, READ)

WORKSTATION\_IS\_UNKNOWN (*station*)

### 18.2.6 WORKSTATION\_REDUCE\_CONNECTION

```

WORKSTATION_REDUCE_CONNECTION (
    station : [ Workstation_designator ],
    status   : Requested_connection_status,
    force    : Boolean
)

```

WORKSTATION\_REDUCE\_CONNECTION reduces the connection status of the workstation *station* to the connection status *status*. If *station* is not supplied, the local workstation is assumed.

If the required change of status is a *disconnection*, i.e. the current status of *station* is CONNECTED and the required status is CLIENT or LOCAL, or the current status is CLIENT and the required status is LOCAL, then *force* has the following effect.

- **true:** The operation is performed whether *station* is busy or not. If the connection status change is from CLIENT to LOCAL, the station ceases to be a client. If the connection status change is from CONNECTED to CLIENT, the station ceases to be a server. If the connection status change is from CONNECTED to LOCAL, the station ceases to be a client or a server.
- **false:** The workstation *station* must not be busy (see 18.1.2).

In other cases *force* has no effect.

#### Errors

ACCESS\_ERRORS (*object*, ATOMIC, CHANGE, CONTROL\_OBJECT)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_CONFIGURATION)

WORKSTATION\_IS\_BUSY (*station*)

WORKSTATION\_IS\_NOT\_CONNECTED (*station*)

WORKSTATION\_IS\_UNKNOWN (*station*)

## 18.3 Foreign system operations

### 18.3.1 CONTENTS\_COPY\_FROM\_FOREIGN\_SYSTEM

```

CONTENTS_COPY_FROM_FOREIGN_SYSTEM (
    file                : File_designator,
    foreign_system      : Foreign_system_designator,
    foreign_name        : String,
    foreign_parameters  : [ String ]
)

```

CONTENTS\_COPY\_FROM\_FOREIGN\_SYSTEM copies the file identified by *foreign\_name*, residing on the foreign system *foreign\_system*, into the file contents of the object *file*, overwriting any previous contents.

The syntax and interpretation of *foreign\_name* and *foreign\_parameters*, whether *foreign\_parameters* is required, and the interpretation of the process's mandatory and discretionary context and the permissions required on *foreign\_name*, are all implementation-defined and may depend on *foreign\_system*.



A write lock of the default mode is obtained on *file*. A read lock of the default mode is obtained on *foreign\_system*.

#### Errors

ACCESS\_ERRORS (*file*, ATOMIC, MODIFY, WRITE\_CONTENTS)

ACCESS\_ERRORS (*foreign\_system*, ATOMIC, READ, NAVIGATE)

FOREIGN\_OBJECT\_IS\_INACCESSIBLE (*foreign\_system*, *foreign\_name*)

FOREIGN\_SYSTEM\_IS\_INACCESSIBLE (*foreign\_system*)

FOREIGN\_SYSTEM\_IS\_UNKNOWN (*foreign\_system*)

STATIC\_CONTEXT\_IS\_IN\_USE (*file*)

### 18.3.2 CONTENTS\_COPY\_TO\_FOREIGN\_SYSTEM

```

CONTENTS_COPY_TO_FOREIGN_SYSTEM (
    file                : File_designator,
    foreign_system      : Foreign_system_designator,
    foreign_name        : String,
    foreign_parameters : [ String ]
)

```

CONTENTS\_COPY\_TO\_FOREIGN\_SYSTEM copies the file contents of the object *object* into a file identified by *foreign\_name* on the foreign system *foreign\_system*.

The syntax and interpretation of *foreign\_name* and *foreign\_parameters*, whether *foreign\_parameters* is required, and the interpretation of the process's mandatory and discretionary context and the permissions required on *foreign\_name*, are all implementation-defined and may depend on *foreign\_system*.

A read lock of the default mode is obtained on *file*. A read lock of the default mode is obtained on *foreign\_system*.

#### Errors

ACCESS\_ERRORS (*file*, ATOMIC, READ, READ\_CONTENTS)

ACCESS\_ERRORS (*foreign\_system*, ATOMIC, READ, NAVIGATE)

FOREIGN\_OBJECT\_IS\_INACCESSIBLE (*foreign\_system*, *foreign\_name*)

FOREIGN\_EXECUTION\_IMAGE\_IS\_BEING\_EXECUTED (*foreign\_system*, *foreign\_name*)

FOREIGN\_SYSTEM\_IS\_INACCESSIBLE (*foreign\_system*)

FOREIGN\_SYSTEM\_IS\_UNKNOWN (*foreign\_system*)

LABEL\_IS\_OUTSIDE\_RANGE (*file*, *foreign\_system*)

NOTE - It is implementation-defined whether the contents of *object* overwrites or is appended to the contents of the foreign file; this may depend on the properties of *foreign\_system* and on *foreign\_parameters*.

## 18.4 Time operations

### 18.4.1 TIME\_GET

```
TIME_GET (  
    )  
    time : Time
```

TIME\_GET returns as *time* the current value of the system time.

#### Errors

None.

### 18.4.2 TIME\_SET

```
TIME_SET (  
    time : Time  
    )
```

TIME\_SET sets the value of system time to *time*.

#### Errors

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_CONFIGURATION)

The following implementation-defined error may be raised:

TIME\_CANNOT\_BE\_CHANGED

## 19 Discretionary security

### 19.1 Discretionary security concepts

#### 19.1.1 Security groups

Group\_identifier = Natural

**sds** discretionary\_security:

**import object type** system-object, system-static\_context, system-process, system-common\_root;

**import attribute type** system-name, system-number;

security\_group\_directory: **child type of object with link**

known\_security\_group: (**navigate**) **existence link** (group\_identifier: **natural**) to security\_group;

security\_groups\_of: **implicit link to common\_root reverse** security\_groups;

**end** security\_group\_directory;

security\_group: **child type of object**;

user: **child type of security\_group with link**

user\_identity\_of: (**navigate**) **non\_duplicated designation link** (number) to process;

user\_member\_of: (**navigate**) **reference link** (number) to user\_group **reverse** has\_users;

**end** user;

user\_group: **child type of security\_group with link**

has\_users: (**navigate**) **reference link** (number) to user **reverse** user\_member\_of;

user\_subgroup\_of: (**navigate**) **reference link** (number) to user\_group **reverse**

has\_user\_subgroups;

has\_user\_subgroups: (**navigate**) **reference link** (number) to user\_group **reverse**

user\_subgroup\_of;

adopted\_user\_group\_of: (**navigate**) **non\_duplicated designation link** (number) to process;

**end** user\_group;

program\_group: **child type of security\_group with link**

has\_programs: (**navigate**) **reference link** (number) to static\_context **reverse**

program\_member\_of;

program\_subgroup\_of: (**navigate**) **reference link** (number) to program\_group **reverse**

has\_program\_subgroups;

has\_program\_subgroups: (**navigate**) **reference link** (number) to program\_group **reverse**

program\_subgroup\_of;

**end** program\_group;

**extend object type** static\_context **with link**

program\_member\_of: (**navigate**) **implicit link** (system\_key) to program\_group **reverse**

has\_programs;

**end** static\_context;

**extend object type process with  
link**

user\_identity: (**navigate**) **designation link to** user;  
 adopted\_user\_group: (**navigate**) **designation link to** user\_group;  
 adoptable\_user\_group: (**navigate**) **designation link (number) to** user\_group **with  
attribute**  
     adoptable\_for\_child: (**read**) **boolean := true;**  
**end adoptable\_user\_group;**  
**end process;**

**extend object type common\_root with  
link**

security\_groups: (**navigate**) **existence link to** security\_group\_directory **reverse**  
     security\_groups\_of;  
**end common\_root;**  
**end discretionary\_security;**

The security group directory is an administrative object (see 9.1.2).

A user is a *member* of a user group if there is a "has\_users" link from the user group to the user.

A static context is a *member* of a program group if there is a "has\_programs" link from the program group to the static context.

A user group A is a *user subgroup* of a user group B if there is a "has\_user\_subgroups" link from the user group B to the user group A. User group B is a *direct user supergroup* of user group A.

An *indirect user supergroup* of a user group is a direct user supergroup of a direct or indirect user supergroup of the user group. A *user supergroup* of a user group is a direct user supergroup or an indirect user supergroup of that user group.

The set of user groups with the user-subgroup/user-supergroup relation forms an acyclic graph with the predefined user group ALL\_USERS as root.

A program group consists of a set of static contexts. A program group A is a *program subgroup* of a program group B if there is a "has\_program\_subgroups" link from the program group B to the program group A. Program group B is a *direct program supergroup* of program group A.

An *indirect program supergroup* of a program group is a direct program supergroup of a direct or indirect program supergroup of the program group. A *program supergroup* of a program group is a direct program supergroup or an indirect program supergroup of that program group.

Where there is no risk of ambiguity, a user subgroup or a program subgroup is called simply a *subgroup*, and a user supergroup or a program supergroup is called simply a *supergroup*.

*Discretionary groups* are security groups used for the purposes of discretionary access control. Each process has the following *effective security groups*:

- One user, the destination of the "user\_identity" link from the process, called *the user* of the process.
- One user group, the *adopted user group* of the process, of which the user is a member, and all user supergroups of that user group (including the group ALL\_USERS). The adopted user group is the destination of the "adopted\_user\_group" link from the process.

- All program groups of which a non-interpretable static context run by a process (see 13.1.1) is a member, and all their supergroups; and for an interpretable static context, the program groups of which the interpreter is a member, and all their supergroups.

Each process also has an associated set of user groups called its *adoptable* user groups which are the destination of "adoptable\_user\_group" links from the process; these are the set of user groups out of which the process may make effective one user group in place of the currently adopted user group. Adoptable user groups must have the user as a member.

When a process creates a child process, its adoptable user groups are inherited except when the "adoptable\_for\_child" attribute of the "adoptable\_user\_group" link from the parent process is **false**.

No object type is a descendant type of more than one of the object types "user", "user\_group", and "program\_group".

The predefined user group ALL\_USERS always exists, as do the predefined program groups PCTE\_SECURITY, PCTE\_AUDIT, PCTE\_EXECUTION, PCTE\_REPLICATION, PCTE\_CONFIGURATION, PCTE\_HISTORY, and PCTE\_SCHEMA\_UPDATE, which are objects in the initial state of the object base linked to the security group directory with predefined values of their group identifiers. Their security group identifiers are as follows:

- |                      |   |
|----------------------|---|
| - ALL_USERS          | 1 |
| - PCTE_SECURITY      | 2 |
| - PCTE_AUDIT         | 3 |
| - PCTE_EXECUTION     | 4 |
| - PCTE_REPLICATION   | 5 |
| - PCTE_CONFIGURATION | 6 |
| - PCTE_HISTORY       | 7 |
| - PCTE_SCHEMA_UPDATE | 8 |

Zero is not used as a security group identifier; it is used to denote absence of a security group.

A user must be a member of a user group in order for a process to act on its behalf.

#### NOTES

1 Discretionary access to objects is controlled on the basis of the effective security groups of the accessing process. Security groups are of three types: users, user groups and program groups. Each process has one group which represents the user on behalf of whom the process is running. A user may play several different roles while using the PCTE-based environment, and these roles are represented by the user groups to which the user belongs. The role the user is playing at any one time is given by the user group which is currently adopted plus its supergroups recursively. It is an important security requirement that a user adopts at most one role before operations are carried out on its behalf. The subgroup structure is intended to reflect the organization of the project into working groups or teams and team membership.

2 Rights may also be granted to a program, which the user also obtains when the program executes on the user's behalf provided that the user has the right to execute the program. Program groups may be used to deny as well as to grant access to specific data objects. In this way program groups may be used to model data abstraction and implement information hiding. They also provide a less specific way of restricting access. A process may only act on behalf of a single user and user group at any one time and which must be authenticated. Giving a right to a program means that the right is given to any user who has the right to execute the program when the program is

executed on behalf of that user. Program groups also provide a means of expanding the number of effective security groups without violating the constraint of there being only one user role effective at any one time.

3 A user which is a member of a user group need not be a member of a sub- or supergroup of that group.

4 The security group structure is intended to be used by tools, such as "login" tools, built on top of PCTE.

5 The predefined user group ALL\_USERS is effective for all processes, as it is the root of the directed acyclic graph of user groups. Access rights which are effective for all users can be given to this user group.

6 A process may have no effective program group.

7 The predefined program groups have the following meanings:

- PCTE\_AUDIT This program group is required by the following operations for the audit mechanism:
  - . AUDIT\_SWITCH\_ON\_SELECTION;
  - . AUDIT\_SWITCH\_OFF\_SELECTION;
  - . AUDIT\_ADD\_CRITERION;
  - . AUDIT\_REMOVE\_CRITERION;
  - . AUDIT\_GET\_CRITERIA;
  - . AUDIT\_SELECTION\_CLEAR;
  - . AUDITING\_STATUS;
  - . AUDIT\_FILE\_COPY\_AND\_RESET.
- PCTE\_CONFIGURATION This program group is required when type identifiers are used to denote invisible types in type references (see 23.1.2.5), and by the following operations which define devices or volumes or which manage workstations or archives:
  - . ARCHIVE\_RESTORE;
  - . ARCHIVE\_SAVE;
  - . DEVICE\_CREATE;
  - . DEVICE\_REMOVE;
  - . VOLUME\_CREATE;
  - . VOLUME\_DELETE;
  - . WORKSTATION\_REDUCE\_CONNECTION;
  - . WORKSTATION\_CREATE;
  - . WORKSTATION\_CREATE;
  - . WORKSTATION\_CONNECT;
  - . WORKSTATION\_DELETE;
- PCTE\_EXECUTION This program group may be required by the following operations for execution mechanisms such as setting the file size limit for a process or changing the priority of a process:
  - . PROCESS\_SET\_FILE\_SIZE\_LIMIT;
  - . PROCESS\_INTERRUPT\_OPERATION;
  - . PROCESS\_SET\_PRIORITY;
  - . TIME\_SET.
- PCTE\_HISTORY. This program group is required by the following operations to explicitly set the last access time or last modification time of an object, or to manipulate the version graph:

- . OBJECT\_SET\_TIME\_ATTRIBUTES;
- . VERSION\_ADD\_PREDECESSOR;
- . VERSION\_REMOVE;
- . VERSION\_REMOVE\_PREDECESSOR.
- PCTE\_REPLICATION This program group is required by the operations of the replication clause and all the operations which modify the object base when they apply to masters of replicated objects. These are a very large subset of all PCTE operations (see C.3). They are not listed here.
- PCTE\_SECURITY This program group is required to use the operations which are critical to either the consistency of the security group structure or to security (or both). These are the three operations:
  - . GROUP\_REMOVE;
  - . GROUP\_RESTORE;
  - . PROCESS\_SET\_USER.
- PCTE\_SCHEMA\_UPDATE. This program group is required by operations which update an SDS, i.e. those defined in 10.2.

### 19.1.2 Access control lists

Discretionary\_access\_mode = APPEND\_CONTENTS | APPEND\_IMPLICIT | APPEND\_LINKS |  
 CONTROL\_DISCRETIONARY | CONTROL\_MANDATORY | CONTROL\_OBJECT |  
 DELETE | EXECUTE | EXPLOIT\_CONSUMER\_IDENTITY | EXPLOIT\_DEVICE |  
 EXPLOIT\_SCHEMA | NAVIGATE | OWNER | READ\_ATTRIBUTES | READ\_CONTENTS |  
 READ\_LINKS | STABILIZE | WRITE\_ATTRIBUTES | WRITE\_CONTENTS |  
 WRITE\_IMPLICIT | WRITE\_LINKS.

Discretionary\_access\_mode\_value = GRANTED | UNDEFINED | DENIED | PARTIALLY\_DENIED

Discretionary\_access\_modes = **set of** Discretionary\_access\_mode

Access\_rights = **map** Discretionary\_access\_mode **to** Discretionary\_access\_mode\_value

Acl = **map** Group\_identifier **to** Access\_rights

Atomic\_discretionary\_access\_mode\_value = GRANTED | UNDEFINED | DENIED

Atomic\_access\_rights = **map** Discretionary\_access\_mode **to**  
 Atomic\_discretionary\_access\_mode\_value

**sds** discretionary\_security:

**import object type** system-object, system-process;

**extend object type** object **with**  
**attribute**

atomic\_acl: **(protected) non\_duplicated string**;  
 composite\_acl: **(protected) non\_duplicated string**;

**end** object;

**extend object type** process **with**  
**attribute**

default\_atomic\_acl: **(protected) string**;  
 default\_object\_owner: **(protected) natural**;

**end** process;

**end** discretionary\_security;



Each object has two associated *access control lists* (or *ACLs*): an *atomic ACL* and a *composite ACL*. They are represented by two string attributes, "atomic\_acl" and "composite\_acl" respectively. The *scope* of an ACL is the set of atomic objects to which it governs access: the scope of the atomic ACL of an object is the atomic object associated with the object; the scope of the composite ACL is the atoms of the object.

Each ACL is a set of ACL entries. Each ACL entry gives the discretionary access mode value of each discretionary access mode for one security group.

In an atomic ACL, the possible discretionary mode values are GRANTED, DENIED, and UNDEFINED. In an composite ACL they are GRANTED, DENIED, UNDEFINED, and PARTIALLY\_DENIED.

*Access right evaluation for a group* is defined by the function

```
EVALUATE_GROUP (
  g : Security_group_designator;
  o : Object_designator;
  s : Object_scope;
  m : Discretionary_access_mode
)
  v : Discretionary_access_mode_value
```

where  $v$  is the discretionary access mode value of  $m$  in the ACL entry for  $g$  in the atomic ACL (if  $s$  is ATOMIC) or the composite ACL (if  $s$  is COMPOSITE) for  $o$ . The group  $g$  is said to have the discretionary access mode  $m$  atomically *granted*, *denied*, or *undefined* to the object  $o$ , if  $s$  is ATOMIC and  $v$  is GRANTED, DENIED, or UNDEFINED respectively; and compositely *granted*, *denied*, *undefined*, or *partially denied* if  $s$  is COMPOSITE and  $v$  is GRANTED, DENIED, UNDEFINED, or PARTIALLY\_DENIED respectively.  $v$  is called the *atomic* or *composite  $m$  value* for  $g$  to  $o$ . If  $v$  is GRANTED,  $g$  is said to have the *atomic* or *composite  $m$  discretionary access right* to  $o$ .

For every object  $o$  there is at least one security group  $g$  for which  $\text{EVALUATE\_GROUP}(g, o, \text{ATOMIC}, \text{CONTROL\_DISCRETIONARY}) = \text{GRANTED}$  and at least once security group  $g'$  for which  $\text{EVALUATE\_GROUP}(g', o, \text{ATOMIC}, \text{CONTROL\_MANDATORY}) = \text{GRANTED}$ .

For every object  $o$  there is at least one security group which has the atomic CONTROL\_DISCRETIONARY right to  $o$ , and at least once security group which has the atomic CONTROL\_MANDATORY right to  $o$ .

The following constraints apply to the composite ACL for an object  $o$  and the atomic ACLs of  $o$  and its components, for any security group  $g$  and for any discretionary access mode  $m$  except OWNER and CONTROL\_DISCRETIONARY:

$\text{EVALUATE\_GROUP}(g, o, \text{COMPOSITE}, m) = \text{GRANTED}$  if and only if  $\text{EVALUATE\_GROUP}(g, o, \text{ATOMIC}, m) = \text{GRANTED}$  and  $\text{EVALUATE\_GROUP}(g, c, \text{ATOMIC}, m) = \text{GRANTED}$  for every component  $c$  of  $o$ .

$\text{EVALUATE\_GROUP}(g, o, \text{COMPOSITE}, m) = \text{DENIED}$  if and only if  $\text{EVALUATE\_GROUP}(g, o, \text{ATOMIC}, m) = \text{DENIED}$  and  $\text{EVALUATE\_GROUP}(g, c, \text{ATOMIC}, m) = \text{DENIED}$  for every component  $c$  of  $o$ .

$\text{EVALUATE\_GROUP}(g, o, \text{COMPOSITE}, m) = \text{PARTIALLY\_DENIED}$  if and only if  $\text{EVALUATE\_GROUP}(g, o, \text{ATOMIC}, m) = \text{DENIED}$  or  $\text{EVALUATE\_GROUP}(g, c, \text{ATOMIC}, m) = \text{DENIED}$  for some component  $c$  of  $o$ , and  $\text{EVALUATE\_GROUP}(g, o, \text{ATOMIC}, m) \neq \text{DENIED}$  and  $\text{EVALUATE\_GROUP}(g, c, \text{ATOMIC}, m) \neq \text{DENIED}$  for some component  $c$  of  $o$ .

ATOMIC,  $m$ )  $\neq$  DENIED or EVALUATE\_GROUP ( $g, c, \text{ATOMIC}, m$ )  $\neq$  DENIED for some component  $c$  of  $o$ .

EVALUATE\_GROUP ( $g, o, \text{COMPOSITE}, m$ ) = UNDEFINED in all other cases.

The following constraints apply to the composite ACL of an object  $o$  and the atomic ACLs of  $o$  and its components, for any security group  $g$  and for the discretionary access modes OWNER and CONTROL\_DISCRETIONARY.

- If EVALUATE\_GROUP ( $g, o, \text{COMPOSITE}, \text{OWNER}$ ) = GRANTED then EVALUATE\_GROUP ( $g, c, \text{COMPOSITE}, \text{OWNER}$ ) = GRANTED for every component  $c$  of  $o$ , EVALUATE\_GROUP ( $g, o, \text{ATOMIC}, \text{CONTROL\_DISCRETIONARY}$ ) = GRANTED, and EVALUATE\_GROUP ( $g, c, \text{ATOMIC}, \text{CONTROL\_DISCRETIONARY}$ ) = GRANTED for every component  $c$  of  $o$ .
- If EVALUATE\_GROUP ( $g, o, \text{COMPOSITE}, \text{OWNER}$ ) = DENIED then EVALUATE\_GROUP ( $g, c, \text{COMPOSITE}, \text{OWNER}$ ) = DENIED for every component  $c$  of  $o$ , EVALUATE\_GROUP ( $g, o, \text{ATOMIC}, \text{CONTROL\_DISCRETIONARY}$ ) = DENIED and EVALUATE\_GROUP ( $g, c, \text{ATOMIC}, \text{CONTROL\_DISCRETIONARY}$ ) = DENIED for every component  $c$  of  $o$ .

Access right evaluation for a process is defined by the function

```
EVALUATE_PROCESS (
    p : Process_designator;
    o : Object_designator;
    s : Object_scope;
    m : Discretionary_access_mode
)
a : Boolean
```

The returned value  $a$  is defined from the ACLs of  $o$  in the following way: EVALUATE\_PROCESS ( $p, o, s, m$ ) = **true** if and only if there is at least one effective group  $g$  of  $p$  for which EVALUATE\_GROUP ( $g, o, s, m$ ) = GRANTED, and for every other effective group  $g'$  of  $p$  EVALUATE\_GROUP ( $g', o, s, m$ ) = GRANTED or EVALUATE\_GROUP ( $g', o, s, m$ ) = UNDEFINED. If  $a$  = **true**,  $p$  is said to have the *atomic* or *composite*  $m$  discretionary access right to  $o$ , according as  $s$  is ATOMIC or COMPOSITE respectively.

The default atomic ACL and default object owner are used to determine the atomic ACLs and composite ACLs of objects created by the process (see 19.1.4).

#### NOTES

- 1 A composite ACL is computable from the atomic ACLs of the object and its components, except for the discretionary access mode OWNER.
- 2 The implementation-defined mapping of access control lists to the string attribute values may economize on space by omitting discretionary access modes with value UNDEFINED, and omitting ACL entries with all values UNDEFINED.
- 3 If OWNER is set to UNDEFINED for an object  $o$  and a group  $g$ , the OWNER values for  $g$  to the components of  $o$ , and the CONTROL\_DISCRETIONARY values for  $g$  to  $o$  and its components, are unchanged.

### 19.1.3 Discretionary access modes

The following list describes the meanings of the discretionary access modes, generally in terms of the classes of operations for which they are *needed atomically* or *compositely* on an object *o*, i.e. for which a necessary precondition is that the calling process has the atomic or composite access right *m*, respectively, to *o*. The exact definitions are given by the occurrences of DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED in the operation descriptions; see 19.2.

- APPEND\_CONTENTS. Needed atomically to append to the contents of an object or to send a message to a message queue.
- APPEND\_IMPLICIT. Needed atomically to create new implicit links of an object.
- APPEND\_LINKS. Needed atomically to create new links, other than implicit links, from an object. (This right is not sufficient to write the non-key attributes of such a link.)
- CONTROL\_DISCRETIONARY. Needed atomically on an object to change its atomic ACL (except CONTROL\_MANDATORY). CONTROL\_DISCRETIONARY occurs only in atomic ACLs.
- CONTROL\_MANDATORY. Needed atomically on an object to change its "confidentiality\_label" and "integrity\_label" attributes and to change the CONTROL\_MANDATORY rights of other groups on that object.
- CONTROL\_OBJECT. Needed atomically on an object to convert it to a descendant type, to move it to another volume, to create or delete "predecessor" and "successor" links to and from the object, and to convert a normal object to a master object or vice versa,. Needed on a message queue to change the number of storage units allowed by the message queue.
- DELETE. Needed compositely on an object to delete the object (i.e. to delete the last composition or existence link to the object). DELETE has no effect in an atomic ACL.
- EXECUTE. Needed atomically on a static context to execute the associated program. EXECUTE has no effect on objects of other types.
- EXPLOIT\_CONSUMER\_IDENTITY. Needed atomically on a consumer group to use it as a consumer identity. EXPLOIT\_CONSUMER\_IDENTITY has no effect on objects of other types.
- EXPLOIT\_DEVICE. Needed atomically on a device supporting volume to mount a volume on the device or to unmount a volume from the device. EXPLOIT\_DEVICE has no effect on objects of other types.
- EXPLOIT\_SCHEMA. Needed atomically on an SDS, a type in SDS, or a type to use it in a working schema or to consult the typing information contained in it. EXPLOIT\_SCHEMA has no effect on objects of other types.
- NAVIGATE. Needed atomically on an object to use a link reference of a link of the object in a pathname (see 23.1.2.2); needed atomically on a foreign system to access a file on that foreign system.
- OWNER. OWNER occurs only in composite ACLs. It is needed to modify the composite ACL of an object, except for implicit modification by modification of the atomic ACL of the object or of a component of the object. This modification right includes the OWNER right for any security group on that object, except that CONTROL\_DISCRETIONARY rights

may apply (see below) if there is no owner of the object. Unlike other discretionary access modes in composite ACLs, modification of OWNER values is not automatic and must be done explicitly using OBJECT\_SET\_ACL\_ENTRY.

An object must always have an atomic ACL such that it is possible that a process could exist with a set of effective groups such that the process has the CONTROL\_DISCRETIONARY discretionary access right to that object and that another or the same process could exist with a set of effective groups such that the process has the CONTROL\_MANDATORY discretionary access right to that object.

An *owner* of an object is a security group with OWNER right to the object. There may be more than one owner of an object.

For changing OWNER discretionary access values the following rules hold:

- . An owner may modify the OWNER discretionary access value to an object for itself, and for another security group except when the other group is also an owner of the object.
- . An owner of an object may modify the OWNER discretionary access value for any group to any component of the object.
- . The OWNER discretionary access value for a group to an object may not be modified if that object is a component of an object to which that group has OWNER granted or denied.
- . If no owner for an object exists, then the OWNER discretionary access value may be modified if OWNER is granted for a group to all components of the object (excluding the object, if it is a component of itself) and CONTROL\_DISCRETIONARY is granted for the group to the object and to all its components.
- . The constraints defined in 19.1.2 must be maintained.

OWNER when used in connection with discretionary security does not have a meaning in the accounting sense.

- READ\_ATTRIBUTES. Needed atomically to read the attribute values of an object and to evaluate a link of an object if the evaluation uses the preferred link type and preferred link key of the object (see 23.1.2.5). For some predefined attributes, e.g. "atomic\_acl", the READ\_ATTRIBUTES right is not needed, if the attribute is retrieved by an operation especially defined to retrieve that attribute.

The READ\_ATTRIBUTES right is not needed to read the attribute values of the links of an object.

- READ\_CONTENTS. Needed atomically to read the contents of an object, to save a message queue or a process address space, to save a message queue, or to peek a message in a message queue.
- READ\_LINKS. Needed atomically to read the attributes of the links of an object, or to scan sets of links of an object.
- STABILIZE. Needed to change the stability of an object, i.e. to create or delete a stabilizing link to it or a compositely stabilizing link to it or to an object of which it is a component.

- WRITE\_ATTRIBUTES. Needed atomically to change the attribute values of an object. It does not control changing the attribute values of the links of an object, nor the time attributes of an object.
- WRITE\_CONTENTS. Needed atomically to write to or update the contents of an object or a process address space, to set or remove a breakpoint in a process, to restore a message queue, or to receive or delete a message from a message queue. An object's contents may not be deleted although it may be emptied.
- WRITE\_IMPLICIT. Needed atomically to delete implicit links of an object. For this category of link, there are no attributes to change.
- WRITE\_LINKS. Needed atomically to delete links, other than implicit links, of an object and to change values of link attributes.

Where EXPLOIT\_SCHEMA, EXPLOIT\_DEVICE, EXPLOIT\_CONSUMER\_IDENTITY, CONTROL\_OBJECT, CONTROL\_DISCRETIONARY, CONTROL\_MANDATORY or OWNER discretionary access rights to an object are required of the calling process by an operation which changes the links or attributes of that object, discretionary access rights which would be appropriate for such changes (e.g. APPEND\_LINKS, WRITE\_ATTRIBUTES) are not also required to that object.

#### NOTES

- 1 OWNER consistency rules demand that an owner may not modify the OWNER discretionary access right to an object for another security group not only when the other security group is an owner of the object, but also when the other security group is the owner of an object of which the object is a component.
- 2 The rules for conferring the OWNER discretionary access right on an object for which no owner exists also cover the case where no owner exists for an object of which the object is a component, since if such an owner existed, an owner would exist for the object under consideration.
- 3 The OWNER right on an object for a security group can never be PARTIALLY\_DENIED. This is achieved by ensuring that when a composition link is created (e.g. by OBJECT\_CREATE, SDS\_CREATE\_OBJECT\_TYPE, or LINK\_RESTORE) any OWNER rights of the newly enclosing object are propagated to the new component, and that when the OWNER right is set on an object (by OBJECT\_SET\_ACL\_ENTRY) the new value is consistent with rights on enclosing objects.

### 19.1.4 Access control lists on object creation

When an object is created, its atomic ACL is determined from the default atomic ACL of the creating process as follows. For each ACL entry in the default atomic ACL, with access rights M, an ACL entry in the atomic ACL is created for the same group with access rights M', where the mapping M' is determined for each discretionary access mode *m* by the corresponding discretionary mode values of M and the access mask A:

- $M'(m) = \text{GRANTED}$  if  $M(m) = \text{GRANTED}$  or  $A(m) = \text{GRANTED}$ , and neither  $M(m) = \text{DENIED}$  nor  $A(m) = \text{DENIED}$ .
- $M'(m) = \text{DENIED}$  if  $M(m) = \text{DENIED}$  or  $A(m) = \text{DENIED}$ .
- $M'(m) = \text{UNDEFINED}$  if  $M(m) = \text{UNDEFINED}$  and  $A(m) = \text{UNDEFINED}$ .

The access mask A is a parameter to the operation used to create the object (e.g. OBJECT\_CREATE).

The default object owner of the creating process defines the group identifier of a security group. The composite ACL of the created object is derived from the atomic ACLs of the object subject to the constraints given in 19.1.2 for all discretionary access modes except OWNER and CONTROL\_DISCRETIONARY. An entry in the composite ACL relates to the default object owner of the creating process, if one exists, and has the OWNER discretionary access mode granted. It is an error if the ACL entry in the created atomic ACL for the default object owner group does not have CONTROL\_DISCRETIONARY granted.

When an operation creates an object, any further accesses to that object during that same operation call are not subject to discretionary or mandatory access checks.

## 19.2 Operations for discretionary access control operation

### 19.2.1 GROUP\_GET\_IDENTIFIER

```
GROUP_GET_IDENTIFIER (
    group      : Security_group_designator
)
    identifier : Group_identifier
```

GROUP\_GET\_IDENTIFIER returns in *identifier* the key of the "known\_security\_group" link from the security group directory to the security group *group*.

#### Errors

ACCESS\_ERRORS (security group determined by *group*, ATOMIC, READ, READ\_LINKS)  
GROUP\_IDENTIFIER\_IS\_INVALID (*group*)

### 19.2.2 OBJECT\_CHECK\_PERMISSION

```
OBJECT_CHECK_PERMISSION (
    object      : Object_designator,
    modes       : Discretionary_access_modes,
    scope       : Object_scope
)
    accessible : Boolean
```

OBJECT\_CHECK\_PERMISSION tests if the calling process has the discretionary and mandatory permission to access the object *object* according to the set of access modes given in *modes* and the scope *scope*. For the discretionary permissions, the operation evaluates EVALUATE\_PROCESS (calling process, *object*, *scope*, *mode*) (see 19.1.2) for each discretionary access mode *mode* in *modes*. For the mandatory permissions read and write access is interpreted according to the discretionary access modes:

- Read access is tested if *modes* contains NAVIGATE, READ\_ATTRIBUTES, READ\_LINKS, READ\_CONTENTS, EXECUTE, EXPLOIT\_DEVICE, EXPLOIT\_SCHEMA or EXPLOIT\_CONSUMER\_IDENTITY.
- Write access is tested if *modes* contains APPEND\_CONTENTS, APPEND\_LINKS, APPEND\_IMPLICIT, WRITE\_ATTRIBUTES, WRITE\_CONTENTS, WRITE\_LINKS, WRITE\_IMPLICIT, DELETE, CONTROL\_DISCRETIONARY, CONTROL\_MANDATORY, CONTROL\_OBJECT or OWNER.



Testing for mandatory read access permission means checking for confidentiality violation and integrity confinement violation (see 20.1). Testing for mandatory write access permission means checking for confidentiality confinement violation and integrity violation. These checks are defined in terms of label domination between the mandatory labels of *object* and the mandatory context of the process.

A read lock of the default mode is obtained on *object*.

The return value *accessible* is:

- **false** if for at least one of the discretionary access modes given in *modes* EVALUATE\_PROCESS (calling process, *object*, *scope*, *mode*) = **false**.
- **false** if read access is implied by *modes* and either LABEL\_DOMINATES (*confidentiality\_context* (process), *confidentiality\_label* (*object*)) = **false** or LABEL\_DOMINATES (*integrity\_label* (*object*), *integrity\_context* (process)) = **false** (see 20.1.3).
- **false** if write access is implied by *modes* and either LABEL\_DOMINATES (*confidentiality\_label* (*object*), *confidentiality\_context*(process)) = **false** or LABEL\_DOMINATES (*integrity\_context*(process), *integrity\_label* (*object*)) = **false**
- **true** otherwise. In this case, for all of the discretionary access modes given in *modes* EVALUATE\_PROCESS (calling process, *object*, *scope*, *mode*) = **true**; and:
  - . if read access is implied by *modes* then LABEL\_DOMINATES (*confidentiality\_context* (process), *confidentiality\_label* (*object*)) = **true** and LABEL\_DOMINATES (*integrity\_label* (*object*), *integrity\_context* (process)) = **true**;
  - . if write access is implied by *modes* then LABEL\_DOMINATES (*confidentiality\_label* (*object*), *confidentiality\_context*(process)) = **true** and LABEL\_DOMINATES (*integrity\_context* (process), *integrity\_label* (*object*)) = **true**.

For the maps *confidentiality\_label*, *confidentiality\_context*, *integrity\_label*, and *integrity\_context* see 20.1.4.

### Errors

ACCESS\_MODE\_IS\_INCOMPATIBLE (*scope*, *modes*)

CONFIDENTIALITY\_WOULD\_BE\_VIOLATED (*granule*, *scope*)

INTEGRITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED (*granule*, *scope*)

OBJECT\_IS\_ARCHIVED (*granule*)

OBJECT\_IS\_INACCESSIBLE (*granule*, *scope*)

NOTE - READ\_ATTRIBUTES access right is not necessary to perform this operation. If it were, the operation would lose much of its usefulness, since access checks do not require any access permissions to read mandatory labels or ACLs.



### 19.2.3 OBJECT\_GET\_ACL

```

OBJECT_GET_ACL (
    object      : Object_designator,
    scope       : Object_scope
)
    acl        : Acl

```

OBJECT\_GET\_ACL returns the atomic or composite ACL of the object *object*, according as *scope* is ATOMIC or COMPOSITE.

A read lock of the default mode is obtained on *object*.

#### Errors

ACCESS\_ERRORS (*granule*, *scope*, READ, READ\_ATTRIBUTES)

NOTE - It is expected that implementations calculate the composite ACL of an object (except for the OWNER modes) from the atomic ACLs of the object and its components.

### 19.2.4 OBJECT\_SET\_ACL\_ENTRY

```

OBJECT_SET_ACL_ENTRY (
    object      : Object_designator,
    group       : Group_identifier,
    modes       : Atomic_access_rights,
    scope       : Object_scope
)

```

OBJECT\_SET\_ACL\_ENTRY sets an ACL entry in the atomic or composite ACL of the object *object* for the security group *group*. If the settings of the ACL entry are already as required, except for setting a composite ACL entry to UNDEFINED, then this operation has no effect. In the case where for *scope* = COMPOSITE, and some mode *m*, *modes*(*m*) is UNDEFINED and EVALUATE\_GROUP (*group*, *object*, COMPOSITE, *m*) is previously UNDEFINED, then there is an effect if for one or more components *c* of *object*, EVALUATE\_GROUP (*group*, *c*, ATOMIC, *m*) is not already set UNDEFINED. EVALUATE\_GROUP (*group*, *c*, ATOMIC, *m*) of such components is changed to UNDEFINED.

If *scope* is ATOMIC, then OBJECT\_SET\_ACL\_ENTRY sets the ACL entry for *group* in the atomic ACL of *object* to *modes*, for all discretionary access modes specified in *modes*. OWNER must not appear in *modes*.

If *scope* is COMPOSITE, then for *object* and all its components, OBJECT\_SET\_ACL\_ENTRY sets the ACL entries for *group* in the composite ACLs and also in the atomic ACLs for all discretionary access modes specified in *modes*, except CONTROL\_DISCRETIONARY and OWNER, to *modes*. CONTROL\_DISCRETIONARY must not appear in *modes*. OWNER is treated as follows:

- the OWNER discretionary access mode value for *group* in the composite ACL of *object* and the CONTROL\_DISCRETIONARY discretionary access mode value for *group* in the atomic ACL of *object* are set to *modes* (OWNER) provided that any outer object of *object* has OWNER undefined for *group*.

- If *modes* (OWNER) = UNDEFINED, then in the components of *object* the discretionary access mode values for OWNER in the composite ACL and the discretionary access mode values for CONTROL\_DISCRETIONARY in the atomic ACL of *group* are not changed.
- If *modes* (OWNER) = GRANTED or DENIED, then in the components of *object* *group* is set to have OWNER granted or denied respectively in the composite ACL, and CONTROL\_DISCRETIONARY granted or denied respectively in the atomic ACL; provided that any outer object of *object* has OWNER undefined for *group*.
- If no owner for *object* exists, then the operation can modify OWNER, and OWNER only, if CONTROL\_DISCRETIONARY is granted for *group* in the atomic ACL of *object*, and for all components of *object*, OWNER is granted for *group* in the composite ACL.

Whether *scope* is ATOMIC or COMPOSITE, the composite ACLs of all outer objects of *object*, and of *object* itself if *scope* is ATOMIC, are updated, so that all constraints defined for composite ACLs, and the atomic and composite ACLs of their components (see 19.1.2) are maintained. OWNER modes of outer objects of *object* are not updated; if this would be necessary to maintain the constraints then an error is raised.

If *scope* is COMPOSITE, then write locks of the default mode are obtained on *object* and on all its components.

### Errors

If *scope* is ATOMIC:

ACCESS\_ERRORS (*object*, ATOMIC, CHANGE, CONTROL\_DISCRETIONARY)

If *scope* is COMPOSITE:

If there is no owner for *object*, and only OWNER is to be modified:

ACCESS\_ERRORS (*object*, ATOMIC, CHANGE, CONTROL\_DISCRETIONARY)

ACCESS\_ERRORS (component of *object*, COMPOSITE, CHANGE, OWNER)

If there is an owner for *object*, or there is no owner for *object* and modes other than OWNER are required to be modified:

ACCESS\_ERRORS (*object*, COMPOSITE, CHANGE, OWNER)

If the CONTROL\_MANDATORY access right is changed:

ACCESS\_ERRORS (*object*, *scope*, CHANGE, CONTROL\_MANDATORY)

If *scope* is COMPOSITE and *modes* (CONTROL\_MANDATORY) = UNDEFINED but no change to the composite ACL is required, then for each atom A of *object* where CONTROL\_MANDATORY is to be changed from GRANTED:

ACCESS\_ERRORS (A, ATOMIC, CHANGE, CONTROL\_MANDATORY)

ACCESS\_MODE\_IS\_NOT\_ALLOWED (*modes*, *scope*)

CONTROL\_WOULD\_NOT\_BE\_GRANTED (*object*)

GROUP\_IDENTIFIER\_IS\_INVALID (*group*)

If *scope* is COMPOSITE:

OBJECT\_HAS\_GROUP\_WHICH\_IS\_ALREADY\_OWNER (*object*, *group*)

OBJECT\_OWNER\_CONSTRAINT\_WOULD\_BE\_VIOLATED (*object*)

The following implementation-dependent error may be raised:

OBJECT\_IS\_INACCESSIBLE (outer object of *object*, ATOMIC)

## NOTES

- 1 If an implementation calculates the composite ACL when retrieving it, it may be so designed that it requires the outer objects to be accessible.
- 2 CONTROL\_DISCRETIONARY rather than READ\_ATTRIBUTES or WRITE\_ATTRIBUTES discretionary access right is required to perform this operation. It would be superfluous to require both.
- 3 If *object* is an SDS which may be in use in a working schema, then any change to its composite ACL only has effect when the SDS is next included in a working schema by PROCESS\_SET\_WORKING\_SCHEMA, PROCESS\_CREATE\_AND\_START, or PROCESS\_START.

### 19.3 Discretionary security administration operations

#### 19.3.1 GROUP\_INITIALIZE

```
GROUP_INITIALIZE (
    group      : User_designator | User_group_designator | Program_group_designator
)
    identifier : Group_identifier
```

GROUP\_INITIALIZE adds the security group *group* to the security group directory. A "known\_security\_group" link is created from the master of the security group directory to *group*. The key of this link is set to a system-generated unique value, which is guaranteed never to be re-used as a security group key and is returned as *identifier*.

Write locks of the default mode are obtained on the created links.

#### Errors

ACCESS\_ERRORS (the security group directory, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*group*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

SECURITY\_GROUP\_IS\_KNOWN (*group*)

## NOTES

- 1 The group identifier, which is the same as the key to the "known\_security\_group" link to the object, may be determined using the GROUP\_GET\_IDENTIFIER operation.
- 2 This operation does not change any copies of the security group directory.

#### 19.3.2 GROUP\_REMOVE

```
GROUP_REMOVE (
    group : User_designator | User_group_designator | Program_group_designator
)
```

GROUP\_REMOVE removes the security group *group* from the set of known groups. The "known\_security\_group" link from the security group directory is deleted. If there are no remaining existence or composition links to *group*, then *group* is also deleted. In this case, the "object\_on\_volume" link to *group* is deleted.

The master of the security group directory is always updated by this operation.

Write locks of the default mode are obtained on the deleted links and object.

**Errors**

ACCESS\_ERRORS (the security group directory, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS(*group*, ATOMIC, MODIFY, WRITE\_IMPLICIT)

If the conditions hold for deletion of the "security\_group" object *group*:

ACCESS\_ERRORS (*group*, COMPOSITE, MODIFY, DELETE)

GROUP\_IDENTIFIER\_IS\_INVALID (*group*)

OBJECT\_HAS\_LINKS\_PREVENTING\_DELETION (*group*)

OBJECT\_IS\_IN\_USE\_FOR\_DELETE (*group*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SECURITY)

SECURITY\_GROUP\_IS\_IN\_USE (*group*)

SECURITY\_GROUP\_IS\_PREDEFINED (*group*)

SECURITY\_GROUP\_IS\_REQUIRED\_BY\_OTHER\_GROUPS (*group*)

NOTE - This operation does not change any copies of the security group directory.

**19.3.3 GROUP\_RESTORE**

```
GROUP_RESTORE (
    group      : User_designator | User_group_designator | Program_group_designator
    identifier  : Group_identifier
)
```

GROUP\_RESTORE adds the security group *group* to the security group directory. A "known\_security\_group" link is created from the master of the security group directory to *group*. The group identifier *identifier* is used as the key for this link. This identifier must be a used group identifier, originally generated when initializing a security group which has since been deleted.

Write locks of the default mode are obtained on the created links.

**Errors**

ACCESS\_ERRORS (the security group directory, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*group*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

GROUP\_IDENTIFIER\_IS\_IN\_USE (*identifier*)

GROUP\_IDENTIFIER\_IS\_INVALID (*identifier*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_SECURITY)

SECURITY\_GROUP\_IS\_KNOWN (*group*)

NOTE - This operation does not change any copies of the security group directory.

**19.3.4 PROGRAM\_GROUP\_ADD\_MEMBER**

```
PROGRAM_GROUP_ADD_MEMBER (
    group      : Program_group_designator,
    program    : Static_context_designator
)
```

PROGRAM\_GROUP\_ADD\_MEMBER adds the program *program* to the program group *group*. A "program\_member\_of" link is created from *program* to *group*, together with a "has\_programs" reverse link. The keys of the created links are implementation-dependent.

Write locks of the default mode are obtained on the created links.

#### Errors

ACCESS\_ERRORS (*group*, ATOMIC, APPEND\_LINKS)

ACCESS\_ERRORS (*program*, ATOMIC, MODIFY, APPEND\_IMPLICIT)

SECURITY\_GROUP\_IS\_UNKNOWN (*group*)

STATIC\_CONTEXT\_IS\_ALREADY\_MEMBER (*program*, *group*)

NOTE - Processes which are current executions of *program* do not receive *group* as an addition to their set of effective security groups.

### 19.3.5 PROGRAM\_GROUP\_ADD\_SUBGROUP

```
PROGRAM_GROUP_ADD_SUBGROUP (
    group      : Program_group_designator,
    subgroup   : Program_group_designator
)
```

PROGRAM\_GROUP\_ADD\_SUBGROUP adds the program group *subgroup* to the program group *group*. A "program\_subgroup\_of" link is created from *subgroup* to *group*, together with a "has\_program\_subgroups" reverse link. The keys of the created links are implementation-dependent (see 23.1.2.5).

Write locks of the default mode are obtained on the created links.

#### Errors

ACCESS\_ERRORS (*group*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*subgroup*, ATOMIC, MODIFY, APPEND\_LINKS)

MASTER\_IS\_INACCESSIBLE (some object of the graph of security groups, ATOMIC)

SECURITY\_GROUP\_ALREADY\_HAS\_THIS\_SUBGROUP (*subgroup*, *group*)

SECURITY\_GROUP\_IS\_IN\_USE (*subgroup*)

SECURITY\_GROUP\_IS\_UNKNOWN (*group*)

SECURITY\_GROUP\_IS\_UNKNOWN (*subgroup*)

SECURITY\_GROUP\_WOULD\_BE\_IN\_INVALID\_GRAPH (*subgroup*, *group*)

### 19.3.6 PROGRAM\_GROUP\_REMOVE\_MEMBER

```
PROGRAM_GROUP_REMOVE_MEMBER (
    group      : Program_group_designator,
    program    : Static_context_designator
)
```

PROGRAM\_GROUP\_REMOVE\_MEMBER removes the static context *program* from the group *group*. The "program\_member\_of" link from *program* to *group* and its "has\_programs" reverse link are deleted.

Write locks of the default mode are obtained on the deleted links.

#### Errors

ACCESS\_ERRORS (*group*, ATOMIC, MODIFY, WRITE\_LINKS)  
 ACCESS\_ERRORS (*program*, ATOMIC, MODIFY, WRITE\_IMPLICIT)  
 SECURITY\_GROUP\_IS\_UNKNOWN (*group*)  
 STATIC\_CONTEXT\_IS\_IN\_USE (*program*)  
 STATIC\_CONTEXT\_IS\_NOT\_MEMBER (*program*, *group*)

### 19.3.7 PROGRAM\_GROUP\_REMOVE\_SUBGROUP

```
PROGRAM_GROUP_REMOVE_SUBGROUP (
    group      : Program_group_designator,
    subgroup   : Program_group_designator
)
```

PROGRAM\_GROUP\_REMOVE\_SUBGROUP removes the program group *subgroup* from the program group *group*. The "program\_subgroup\_of" link from *subgroup* to *group* and its "has\_program\_subgroups" reverse link are deleted.

Write locks of the default mode are obtained on the deleted links.

#### Errors

ACCESS\_ERRORS (*group*, ATOMIC, MODIFY, WRITE\_LINKS)  
 ACCESS\_ERRORS (*subgroup*, ATOMIC, MODIFY, WRITE\_LINKS)  
 PROGRAM\_GROUP\_IS\_NOT\_EMPTY (*subgroup*)  
 SECURITY\_GROUP\_IS\_IN\_USE (*subgroup*)  
 SECURITY\_GROUP\_IS\_NOT\_A\_SUBGROUP (*subgroup*, *group*)  
 SECURITY\_GROUP\_IS\_UNKNOWN (*group*)  
 SECURITY\_GROUP\_IS\_UNKNOWN (*subgroup*)

### 19.3.8 USER\_GROUP\_ADD\_MEMBER

```
USER_GROUP_ADD_MEMBER (
    group      : User_group_designator,
    user       : User_designator
)
```

USER\_GROUP\_ADD\_MEMBER adds the user *user* to the user group *group*. A "user\_member\_of" link from *user* to *group* and a "has\_users" reverse link are created. The keys of the created links are implementation-dependent.

Write locks of the default mode are obtained on the created links.

#### Errors

ACCESS\_ERRORS (*group*, ATOMIC, MODIFY, APPEND\_LINKS)  
 ACCESS\_ERRORS (*user*, ATOMIC, MODIFY, APPEND\_LINKS)  
 SECURITY\_GROUP\_IS\_UNKNOWN (*group*)

SECURITY\_GROUP\_IS\_UNKNOWN (*user*)

USER\_GROUP\_LACKS\_ALL\_USERS\_AS\_SUPERGROUP (*group*)

USER\_IS\_ALREADY\_MEMBER (*user*, *group*)

NOTE - This operation does not cause *group* to become an adoptable group of a process running on behalf of *user*.

### 19.3.9 USER\_GROUP\_ADD\_SUBGROUP

```
USER_GROUP_ADD_SUBGROUP (
    group      : User_group_designator,
    subgroup   : User_group_designator
)
```

USER\_GROUP\_ADD\_SUBGROUP adds the user group *subgroup* to the user group *group*. A "user\_subgroup\_of" link from *subgroup* to *group* and a "has\_user\_subgroups" reverse link are created. The keys of the created links are implementation-dependent.

Write locks of the default mode are obtained on the created links.

#### Errors

ACCESS\_ERRORS (*group*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*subgroup*, ATOMIC, MODIFY, APPEND\_LINKS)

MASTER\_IS\_INACCESSIBLE (some object of the graph of security groups, ATOMIC)

SECURITY\_GROUP\_ALREADY\_HAS\_THIS\_SUBGROUP (*subgroup*, *group*)

SECURITY\_GROUP\_IS\_IN\_USE (*subgroup*)

SECURITY\_GROUP\_IS\_UNKNOWN (*group*)

SECURITY\_GROUP\_IS\_UNKNOWN (*subgroup*)

SECURITY\_GROUP\_WOULD\_BE\_IN\_INVALID\_GRAPH (*subgroup*, *group*)

### 19.3.10 USER\_GROUP\_REMOVE\_MEMBER

```
USER_GROUP_REMOVE_MEMBER (
    group : User_group_designator,
    user  : User_designator
)
```

USER\_GROUP\_REMOVE\_MEMBER removes the user *user* from the group *group*. The "user\_member\_of" link from *user* to *group* and its "has\_users" reverse link are deleted.

Write locks of the default mode are obtained on the deleted links.

#### Errors

ACCESS\_ERRORS (*group*, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (*user*, ATOMIC, MODIFY, WRITE\_LINKS)

SECURITY\_GROUP\_IS\_UNKNOWN (*group*)

SECURITY\_GROUP\_IS\_UNKNOWN (*user*)

USER\_GROUP\_IS\_IN\_USE (*user*, *group*)

USER\_IS\_NOT\_MEMBER (*user*, *group*)



NOTE - The "adoptable\_user\_group" link from a process executed on behalf of *user* to *group* is not deleted (see PROCESS\_ADOPT\_USER\_GROUP).

### 19.3.11 USER\_GROUP\_REMOVE\_SUBGROUP

```

USER_GROUP_REMOVE_SUBGROUP (
    group      : User_group_designator,
    subgroup   : User_group_designator
)

```

USER\_GROUP\_REMOVE\_SUBGROUP removes the user group *subgroup* from the user group *group*. The "user\_subgroup\_of" link from *subgroup* to *group* and its "has\_user\_subgroups" reverse link are deleted.

*subgroup* must not be the effective group for a running process.

Write locks of the default mode are obtained on the deleted links.

#### Errors

ACCESS\_ERRORS (*group*, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (*subgroup*, ATOMIC, MODIFY, WRITE\_LINKS)

SECURITY\_GROUP\_IS\_IN\_USE (*subgroup*)

SECURITY\_GROUP\_IS\_NOT\_A\_SUBGROUP (*subgroup*, *group*)

SECURITY\_GROUP\_IS\_UNKNOWN (*group*)

SECURITY\_GROUP\_IS\_UNKNOWN (*subgroup*)

USER\_GROUP\_WOULD\_NOT\_HAVE\_ALL\_USERS\_AS\_SUPERGROUP (*subgroup*)

## 20 Mandatory security

### 20.1 Mandatory security concepts

#### 20.1.1 Mandatory classes

```

sds mandatory_security:
    import object type system-object, system-volume, system-device, system-common_root;
    import attribute type system-name, system-number;
    import object type discretionary_security-security_group, discretionary_security-user;
    import attribute type discretionary_security-group_identifier;
    extend object type security_group with
    link
        may_downgrade: (navigate) reference link (name) to confidentiality_class reverse
            downgradable_by;
        may_upgrade: (navigate) reference link (name) to integrity_class reverse
            upgradable_by;
    end security_group;

```

```

extend object type user with
link
    cleared_for: (navigate) reference link (name) to mandatory_class reverse
        having_clearance;
end user;

mandatory_directory: child type of object with
link
    known_mandatory_class: (navigate) existence link (name) to mandatory_class;
    mandatory_classes_of: implicit link to common_root reverse mandatory_classes;
end mandatory_directory;

mandatory_class: child type of object with
link
    having_clearance: (navigate) reference link (group_identifier) to user reverse
        cleared_for;
end mandatory_class;

confidentiality_class: child type of mandatory_class with
link
    dominates_in_confidentiality: (navigate) reference link to confidentiality_class reverse
        confidentiality_dominator;
    confidentiality_dominator: (navigate) reference link to confidentiality_class reverse
        dominates_in_confidentiality;
    downgradable_by: (navigate) reference link (group_identifier) to security_group
        reverse may_downgrade;
end confidentiality_class;

integrity_class: child type of mandatory_class with
link
    dominates_in_integrity: (navigate) reference link to integrity_class reverse
        integrity_dominator;
    integrity_dominator: (navigate) reference link to integrity_class reverse
        dominates_in_integrity;
    upgradable_by: (navigate) reference link (group_identifier) to security_group reverse
        may_upgrade;
end integrity_class;

extend object type object with
attribute
    confidentiality_label: (read) string;
    integrity_label: (read) string;
end object;

extend object type common_root with
link
    mandatory_classes: (navigate) existence link to mandatory_directory reverse
        mandatory_classes_of;
end common_root;

end mandatory_security;

```

The "mandatory\_class" object type represents the mandatory classes defined for the PCTE installation. The name of a class, the *class name*, is the key attribute of the "known\_mandatory\_class" link from the mandatory directory to the mandatory class object. The destinations of the "having\_clearance" links represent users which are *cleared* to this mandatory class. The concept of user clearance is elaborated in 20.1.4.

The "confidentiality\_class" object type represents the subset of mandatory classes which are confidentiality classes:

- the destination of the "dominates\_in\_confidentiality" link represents the confidentiality class which this confidentiality class dominates;
- the destination of the "confidentiality\_dominator" link represents the confidentiality class which dominates this confidentiality class;
- the destinations of the "downgradable\_by" links represent the security groups which have authority to downgrade from that confidentiality class (see 20.1.4).

The "integrity\_class" object type represents the subset of mandatory classes which are integrity classes:

- the destination of the "dominates\_in\_integrity" link represents the integrity class which this integrity class dominates;
- the destination of the "integrity\_dominator" link represents the integrity class which dominates this integrity class;
- the destinations of the "upgradable\_by" links represent the security groups which have authority to upgrade from that integrity class.

The mandatory directory is an administrative object (see 9.1.2).

### 20.1.2 The mandatory class structure

Confidentiality\_tower = seq1 of Confidentiality\_class\_designator

Integrity\_tower = seq1 of Integrity\_class\_designator

Each mandatory class participates in exactly one of the confidentiality towers and integrity towers which define the dominance relationships between these classes. No class may appear more than once in a tower.

The "dominates\_in\_confidentiality" and "confidentiality\_dominator" links of a mandatory class represent the sequence of confidentiality classes in a confidentiality tower. The destination of a "dominates\_in\_confidentiality" link is the member of the sequence which immediately precedes the origin of the link. The destination of "confidentiality\_dominator" is the member of the sequence which is the immediate successor of the origin of the link.

The "dominates\_in\_integrity" and "integrity\_dominator" links represent the sequence of integrity classes in an integrity tower. The destination of a "dominates\_in\_integrity" link is the member of the sequence which immediately precedes the origin of the link. The destination of "integrity\_dominator" is the member of the sequence which is the immediate successor of the origin of the link.

The predicates CLASS\_DOMINATES and CLASS\_STRICTLY\_DOMINATES are defined in terms of the relative positions of the mandatory classes within a confidentiality tower or an integrity tower. A class *left\_class* dominates a class *right\_class* if CLASS\_DOMINATES (*left\_class*, *right\_class*) = **true**. A class *left\_class* strictly dominates a class *right\_class* if CLASS\_STRICTLY\_DOMINATES (*left\_class*, *right\_class*) = **true**.

```

CLASS_DOMINATES (
    left_class : Mandatory_class_designator,
    right_class : Mandatory_class_designator
)
    result      : Boolean

```

The result is **false** if *left\_class* and *right\_class* do not occur in the same confidentiality tower or integrity tower.

If *left\_class* and *right\_class* occur in a confidentiality tower or an integrity tower *T*, and *left\_class* = *T*(*I*) and *right\_class* = *T*(*J*), then the result is **true** if  $I \geq J$ , otherwise **false**.

```

CLASS_STRICTLY_DOMINATES (
    left_class : Mandatory_class_designator,
    right_class : Mandatory_class_designator
)
    result      : Boolean

```

The result is **false** if *left\_class* and *right\_class* do not occur in the same confidentiality tower or integrity tower.

If *left\_class* and *right\_class* occur in a confidentiality tower or an integrity tower *T*, and *left\_class* = *T*(*I*) and *right\_class* = *T*(*J*), then the result is **true** if  $I > J$ , otherwise **false**.

### 20.1.3 Labels and the concept of dominance

Security\_label = [ Mandatory\_class\_designator ] | Conjunction | Disjunction

Conjunction :: UNITS: **set of** Security\_label

Disjunction :: UNITS: **set of** Security\_label

A security label is either a confidentiality label or an integrity label; the structure is the same in either case.

A class name is a confidentiality or integrity class name. Confidentiality class names may occur only in confidentiality labels. Integrity class names may occur only in integrity labels. Conjunctions and disjunctions must contain at least 2 units. A security label of the first kind in which the optional unit is not supplied is called a *null label*.

The concept of mandatory security permissions depends on the concept of a label dominating another.

The predicates LABEL\_DOMINATES and LABEL\_STRICTLY\_DOMINATES are defined in terms of the possible forms of the labels and the domination relationships between the mandatory classes. A label *left\_label* dominates a label *right\_label* if LABEL\_DOMINATES (*left\_label*, *right\_label*) = **true**. A label *left\_label* strictly dominates a label *right\_label* if LABEL\_STRICTLY\_DOMINATES (*left\_label*, *right\_label*) = **true**.

```

LABEL_DOMINATES (
    left_label : Security_label,
    right_label : Security_label
)
    dominates : Boolean

```

If *right\_label* is null then *dominates* = **true**.

If *left\_label* is null and *right\_label* is not null then *dominates* = **false**.

If *left\_label* and *right\_label* are both mandatory class designators then *dominates* = CLASS\_DOMINATES (*left\_label*, *right\_label*).

If *left\_label* is a mandatory class designator and *right\_label* is a disjunction of mandatory class designators  $r_1, r_2, \dots$  then *dominates* = **true** if CLASS\_DOMINATES (*left\_label*,  $r_j$ ) is **true** for some  $r_j$ , and **false** otherwise.

If *left\_label* is a conjunction of mandatory class designators  $l_1, l_2, \dots$  and *right\_label* is a mandatory class designator then *dominates* = **true** if CLASS\_DOMINATES ( $l_i$ , *right\_label*) is **true** for some  $l_i$ , and **false** otherwise.

If *left\_label* is a conjunction of mandatory class designators  $l_1, l_2, \dots$  and *right\_label* is a disjunction of mandatory class designators  $r_1, r_2, \dots$  then *dominates* = **true** if CLASS\_DOMINATES( $l_i$ ,  $r_j$ ) is **true** for some  $l_i$  and  $r_j$ , and **false** otherwise.

If *right\_label* is a disjunction of security labels  $r_1, r_2, \dots$  and some  $r_k$  is a disjunction of security labels  $r_1', r_2', \dots$  then *dominates* = LABEL\_DOMINATES (*left\_label*,  $r'$ ) where  $r'$  is the disjunction of all the  $r_j'$  and all the  $r_i$  except  $r_k$ .

If *left\_label* is a conjunction of security labels  $l_1, l_2, \dots$  and some  $l_k$  is a conjunction of security labels  $l_1', l_2', \dots$  then *dominates* = LABEL\_DOMINATES ( $l'$ , *right\_label*) where  $l'$  is the conjunction of all the  $l_j'$  and all the  $l_i$  except  $l_k$ .

If *right\_label* is a conjunction of security labels  $r_1, r_2, \dots$  then *dominates* = **true** if LABEL\_DOMINATES (*left\_label*,  $r_j$ ) is **true** for all  $r_j$ , and **false** otherwise.

If *left\_label* is a disjunction of security labels  $l_1, l_2, \dots$  then *dominates* = **true** if LABEL\_DOMINATES ( $l_i$ , *right\_label*) is **true** for all  $l_i$ , and **false** otherwise.

If *right\_label* is a disjunction of security labels  $r_1, r_2, \dots$  and some  $r_k$  is a conjunction of security labels  $r_1', r_2', \dots$  then *dominates* = **true** if LABEL\_DOMINATES (*left\_label*,  $r'$ ) is **true** for all  $r_j'$ , where  $r'$  is *right\_label* with  $r_k$  replaced by  $r_j'$ , and **false** otherwise.

If *left\_label* is a conjunction of security labels  $l_1, l_2, \dots$  and some  $l_k$  is a disjunction of security labels  $l_1', l_2', \dots$  then *dominates* = **true** if LABEL\_DOMINATES ( $l'$ , *right\_label*) is **true** for all  $l_i'$  where  $l'$  is *left\_label* with  $l_k$  replaced by  $l_i'$ , and **false** otherwise.

```

LABEL_STRICTLY_DOMINATES (
    left_label      : Security_label,
    right_label     : Security_label
)
    dominates      : Boolean

```

The definition of this predicate is the same as for LABEL\_DOMINATES except that:

- If *left\_label* and *right\_label* are both null, then *dominates* is **false**.
- CLASS\_DOMINATES is replaced by CLASS\_STRICTLY\_DOMINATES.
- LABEL\_DOMINATES is replaced by LABEL\_STRICTLY\_DOMINATES.

#### NOTES

1 It is possible for label A to dominate label B and B to dominate A, and for label C not to dominate label D while D does not dominate C.

2 For the mapping of security labels to language bindings see 23.1.3.1.

#### 20.1.4 Mandatory rules for information flow

A *user's confidentiality clearance* is a security label derived from the confidentiality classes to which that user is cleared by forming a conjunction of the confidentiality class names.

A *user's integrity clearance* is a security label derived from the integrity classes to which that user is cleared by forming a conjunction of the integrity class names.

A process has a *mandatory context* associated with it which is used to control the flow of information to and from the process. This mandatory context consists of a confidentiality component called a *confidentiality context* and an integrity component called an *integrity context*.

The confidentiality context and integrity context are represented by the "confidentiality\_label" and "integrity\_label" attributes respectively of the process as inherited from the parent type "object" (see 20.1.1).

A process may change its confidentiality context during execution so that the new confidentiality context dominates the previous value.

A process may change its integrity context during execution so that the new integrity context is dominated by the previous value.

A process may only change its confidentiality context so that the result is dominated by the user's confidentiality clearance.

When a confidentiality context is changed, it must remain within the confidentiality label range of the workstation on which the process is executing. When an integrity context is changed, it must remain within the integrity label range of the workstation on which the process is executing (see 20.1.5).

An object has a confidentiality label and an integrity label which control the flow of information into and out of its associated atomic object. The following rules apply to a process' mandatory context or an object's mandatory labels after any change due to the floating of labels (see 20.1.6).

For the purposes of these rules *read* and *write* are defined as follows in terms of information flow. If information flows from an object to a process, the process reads from the object. If information flows from a process to an object, even if it is only erasure, the process writes to the object. Reading and writing refer to any property of an object (attributes, links, link attributes, contents) which can contain (or embody) information. Deletion of an object is therefore considered writing to the object, although deletion of an object is only achieved by deleting a link.

- The simple confidentiality rule: A process P may only read from the atomic object associated with an object A if LABEL\_DOMINATES (confidentiality context of P, confidentiality label of A).
- The confidentiality confinement rule: A process P may only write to the atomic object associated with an object A if LABEL\_DOMINATES (confidentiality label of A, confidentiality context of P).

- The simple integrity rule: A process P may only write to the atomic object associated with an object A if LABEL\_DOMINATES (integrity context of P, integrity label of A).
- The integrity confinement rule: A process P may only read from the atomic object associated with an object A if LABEL\_DOMINATES (integrity label of A, integrity context of P).
- The communication rule: A process P may transmit information to another process Q (by PROCESS\_PEEK or PROCESS\_POKE) if LABEL\_DOMINATES (confidentiality context of Q, confidentiality context of P) and LABEL\_DOMINATES (integrity context of P, integrity context of Q).

A process may change the confidentiality and integrity labels of an object if and only if it has the atomic CONTROL\_MANDATORY right to that object. Under this condition, a confidentiality label may be changed to a value which dominates the previous value and an integrity label may be changed to a value which is dominated by the previous value.

When a confidentiality label of an object is changed, it must remain within the confidentiality label range of the volume on which the object is residing. When an integrity label of an object is changed, it must remain within the integrity label range of the volume on which the object is residing (see 20.1.5). This is true even with the downgrade or upgrade privileges, described below, effective.

If an effective security group of the calling process has additional downgrade or upgrade privileges, these object mandatory labels may be changed so that the new value of the confidentiality label does not dominate the previous value and the new value of the integrity label is not dominated by the previous value, according to the rules defined below:

A process is defined to be *acting with downgrade authority from a confidentiality class C* if the process has an effective security group which has downgrade authority from C, i.e. there is a "downgradable\_by" link from C to the security group. This is represented by the predicate DOWNGRADE\_AUTHORITY:

```
DOWNGRADE_AUTHORITY (
    process : Process_designator,
    class   : Confidentiality_class_designator
)
    authority : Boolean
```

A process is defined to be *acting with upgrade authority to an integrity class C* if the process has an effective group which has upgrade authority to C, i.e. there is an "upgradable\_by" link from C to the security group. This is represented by the predicate UPGRADE\_AUTHORITY:

```
UPGRADE_AUTHORITY (
    process : Process_designator,
    class   : Integrity_class_designator
)
    authority : Boolean
```

A process is permitted to change a confidentiality label from *right\_label* to *left\_label* providing that *right\_label* is *dominated in confidentiality by left\_label relative to the process*. This concept is defined by the following predicates for classes and labels:



```

RELATIVE_CLASS_DOMINATES_IN_CONFIDENTIALITY (
    process : Process_designator,
    left_class : Mandatory_class_designator,
    right_class : Mandatory_class_designator
)
    dominates : Boolean

```

This is the same as CLASS\_DOMINATES except that if DOWNGRADE\_AUTHORITY (process, right\_class) is **true**, dominates is always **true**.

```

RELATIVE_LABEL_DOMINATES_IN_CONFIDENTIALITY (
    process : Process_designator,
    left_label : Security_label,
    right_label : Security_label
)
    dominates : Boolean

```

This is the same as LABEL\_DOMINATES except that:

- The rule beginning 'If left\_label is null' (i.e. the second rule) is replaced by the rule: If left\_label is null and right\_label is a class name C, then dominates = DOWNGRADE\_AUTHORITY (process, C).
- RELATIVE\_CLASS\_DOMINATES\_IN\_CONFIDENTIALITY replaces CLASS\_DOMINATES.
- RELATIVE\_LABEL\_DOMINATES\_IN\_CONFIDENTIALITY replaces LABEL\_DOMINATES.

A process is permitted to change an integrity label from left\_label to right\_label providing that left\_label dominates right\_label in integrity relative to the process. This concept is defined by the following predicates for classes and labels:

```

RELATIVE_CLASS_DOMINATES_IN_INTEGRITY (
    process : Process_designator,
    left_class : Mandatory_class_designator,
    right_class : Mandatory_class_designator
)
    dominates : Boolean

```

This is the same as CLASS\_DOMINATES except that if UPGRADE\_AUTHORITY (process, right\_class) is **true**, dominates is always **true**.

```

RELATIVE_LABEL_DOMINATES_IN_INTEGRITY (
    process : Process_designator,
    left_label : Security_label,
    right_label : Security_label
)
    dominates : Boolean

```

This is the same as LABEL\_DOMINATES except that:

- The rule beginning 'If left\_label is null' (i.e. the second rule) is replaced by the rule: If left\_label is null and right\_label is a class name C, then dominates = UPGRADE\_AUTHORITY (process, C).
- RELATIVE\_CLASS\_DOMINATES\_IN\_INTEGRITY replaces CLASS\_DOMINATES.
- RELATIVE\_LABEL\_DOMINATES\_IN\_INTEGRITY replaces LABEL\_DOMINATES.

The confidentiality context of a process is always dominated by the user's confidentiality clearance, and the integrity clearance of a process is always dominated by the user's integrity clearance.

#### NOTES

- 1 Read and write for mandatory access control are defined in the operations in terms of information flow. If information flows from an object to the process (i.e. access errors may occur with permission READ), it is a read. If information flows from the process to an object (i.e. access errors may occur with permission CHANGE or MODIFY), even if it is only erasure, it is a write. Reading and writing refer to any property of an object (attributes, links, link attributes, contents) which can contain (or embody) information. Deletion of an object is therefore considered a write, although for PCTE, deletion of an object is only achieved by deleting a link.
- 2 A restriction on a process's integrity context with reference to the user's integrity clearance is unnecessary because a change is always a downgrade.
- 3 The restrictions to changes to a process's confidentiality context or integrity context above apply to the operations `PROCESS_SET_CONFIDENTIALITY_LABEL` and `PROCESS_SET_INTEGRITY_LABEL`, and to the floating security labels facility (see 20.1.6) when objects, whose labels would normally prevent access, are read by the process.
- 4 The restrictions to changes to an object's confidentiality or integrity labels above apply to the operations `OBJECT_SET_CONFIDENTIALITY_LABEL` and `OBJECT_SET_INTEGRITY_LABEL`, and to the floating security labels facility when objects, whose labels would normally prevent access, are written to by a process.
- 5 The predicates `RELATIVE_LABEL_DOMINATES_IN_CONFIDENTIALITY` and `RELATIVE_LABEL_DOMINATES_IN_INTEGRITY` are used in operations `OBJECT_SET_CONFIDENTIALITY_LABEL` and `OBJECT_SET_INTEGRITY_LABEL` to define some of these checks.
- 6 In specifying which accesses are read and which write for mandatory access control, the intention is that the rules should be as follows.
  - Each object, its attributes, its contents, its outgoing links (except system-managed designation links representing the use of the object by a process and those representing locks) and their attributes, and its preferred link type and key may be treated as a separate security object.
  - Every access to one of those security objects that depends on data from it may be treated as a read, except that the audit selection criteria are accessible, for the purposes of determining whether the event is auditable, without a mandatory read check, and reading security labels for mandatory access checks does not count as a read.
  - Every access to one of those security objects that writes data to it is treated as a write, with the following exceptions:
    - . the `last_access_time` attribute shall be updatable without mandatory write checks;
    - . records shall be written to the audit file and accounting log without write mandatory checks;
    - . updates arising as a result of process failure or abnormal closedown of a workstation shall be possible without mandatory checks.

### 20.1.5 Multi-level security labels

`Multi_level_device_designator` = `Volume_designator` | `Device_designator` |  
`Execution_site_designator`

**sds** mandatory\_security:

**import object type** system-volume, system-device, system-execution\_site;

**extend object type volume with  
attribute**

confidentiality\_high\_label: (read) non\_duplicated string;  
confidentiality\_low\_label: (read) non\_duplicated string;  
integrity\_high\_label: (read) non\_duplicated string;  
integrity\_low\_label: (read) non\_duplicated string;

**end** volume;

**extend object type device with  
attribute**

confidentiality\_high\_label;  
confidentiality\_low\_label;  
integrity\_high\_label;  
integrity\_low\_label;  
contents\_confidentiality\_label: (read) non\_duplicated string;  
contents\_integrity\_label: (read) non\_duplicated string;

**end** device;

**extend object type execution\_site with  
attribute**

confidentiality\_high\_label;  
confidentiality\_low\_label;  
integrity\_high\_label;  
integrity\_low\_label;

**end** execution\_site;

**end** mandatory\_security;

*Multi-level secure devices* are volumes, devices, and execution sites; they allow data with a fixed range of mandatory labels for confidentiality and for integrity to be stored on them. The fixed ranges of labels required for a multi-level secure device are expressed as two labels, a high label and a low label. In each range the high label must dominate the low label.

A MAXIMUM\_LABEL high end of range means that there is no ceiling on the labels of objects contained within the device.

For it to be permissible for an object *A* to be stored on a multi-level secure device *M*, CONFIDENTIALITY\_LABEL\_WITHIN\_RANGE (*A*, *M*) and INTEGRITY\_LABEL\_WITHIN\_RANGE(*A*, *M*) must be **true**, where:

```
CONFIDENTIALITY_LABEL_WITHIN_RANGE (
  object      : Object_designator,
  device      : Multi_level_device_designator
)
inside_range : Boolean
```

*inside\_range* is **true** if the confidentiality low label of *device* does not strictly dominate the confidentiality label of *object*, and the confidentiality high label of *device* either is MAXIMUM\_LABEL or dominates the confidentiality label of *object*; and is otherwise **false**.

```
INTEGRITY_LABEL_WITHIN_RANGE (
  object      : Object_designator,
  device      : Multi_level_device_designator
)
inside_range : Boolean
```

*inside\_range* is **true** if the integrity low label of *device* does not strictly dominate the integrity label of *object*, and the integrity high label of *device* either is MAXIMUM\_LABEL or dominates the integrity label of *object*; and is otherwise **false**.

Similar checks are made when multi-level secure devices are put on other multi-level secure devices. For it to be permissible for a multi-level secure device A to reside on another multi-level secure device B, CONFIDENTIALITY\_RANGE\_WITHIN\_RANGE(A, B) must be true, where:

```
CONFIDENTIALITY_RANGE_WITHIN_RANGE (
    inner_device : Multi_level_device_designator,
    outer_device : Multi_level_device_designator
)
    inside_range : Boolean
```

*inside\_range* is **true** if the confidentiality low label of *outer\_device* does not strictly dominate the confidentiality low label of *inner\_device*, and the confidentiality high label of *outer\_device* either is MAXIMUM\_LABEL or dominates the confidentiality high label of *inner\_device*; and is otherwise **false**.

```
INTEGRITY_RANGE_WITHIN_RANGE (
    inner_device : Multi_level_device_designator,
    outer_device : Multi_level_device_designator
)
    inside_range : Boolean
```

*inside\_range* is **true** if the integrity low label of *outer\_device* does not strictly dominate the integrity low label of *inner\_device*, and the integrity high label of *outer\_device* either is MAXIMUM\_LABEL or dominates the integrity high label of *inner\_device*; and is otherwise **false**.

The confidentiality or integrity label of an object A *lies within the confidentiality or integrity range* of a multi-level secure device B if CONFIDENTIALITY\_LABEL\_WITHIN\_RANGE(A, B) or INTEGRITY\_LABEL\_WITHIN\_RANGE(A, B) respectively is **true**.

The confidentiality or integrity range of a multi-level secure device A *lies within the confidentiality or integrity range* of a multi-level secure device B if CONFIDENTIALITY\_RANGE\_WITHIN\_RANGE(A, B) or INTEGRITY\_RANGE\_WITHIN\_RANGE(A, B) respectively is **true**.

In addition to its mandatory labels, a device object is associated with two other labels (one confidentiality label and one integrity label), termed *labels of contents*, which govern access to its contents through the device contents operations (see clause 12).

The labels of contents are evaluated in accordance with the characteristics of the physical device each time the contents of the device is accessed. If the device is open for reading or writing, the label associated with the contents of the physical device is implementation-defined. If the label cannot be identified then the confidentiality label C of the contents is set to the confidentiality context of the accessing process, and the integrity label I of the contents is set to the integrity context of the accessing process. If *device\_contents* is defined as a pseudo-object representing the contents which have the labels of contents, the process is denied access to the contents if any of the following are **false**:

- CONFIDENTIALITY\_LABEL\_WITHIN\_RANGE(*device\_contents*, *device*)

- INTEGRITY\_LABEL\_WITHIN\_RANGE (*device\_contents*, *device*)
- LABEL\_DOMINATES (confidentiality context of *process*, C)
- LABEL\_DOMINATES (I, integrity context of *process*)

## NOTES

1 Checks are made that the constraints are obeyed on an object stored on a multi-level secure device whenever:

- objects have their labels changed;
- processes have their mandatory context changed;
- objects are put on to multi-level secure devices:
  - . objects are created on a volume;
  - . objects are moved to a volume;
  - . processes are started or called on a workstation;
- copying files to foreign system.

2 The checks made that the constraints are obeyed when multi-level secure devices are put on other multi-level secure devices apply in the following situations:

- volumes are created on devices;
- volumes are mounted on devices;
- devices are created on workstations;
- security ranges on multi-level devices are changed (see 20.2.9 and 20.2.10).

### 20.1.6 Floating security levels

Floating\_level = NO\_FLOAT | FLOAT\_IN | FLOAT\_OUT | FLOAT\_IN\_OUT

**sds** mandatory\_security:

**import object type** system-process;

floating\_level: NO\_FLOAT, FLOAT\_IN, FLOAT\_OUT, FLOAT\_IN\_OUT;

**extend object type** process **with**  
**attribute**

floating\_confidentiality\_level: (read) non\_duplicated enumeration (floating\_level) :=  
NO\_FLOAT;

floating\_integrity\_level: (read) non\_duplicated enumeration (floating\_level) :=  
NO\_FLOAT;

**end** process;

**end** mandatory\_security;

The floating security levels mechanism enables a process to select either or both of the two facilities:

- The mandatory context of a process may float up (confidentiality) or down (integrity) when information is read from an object.
- The mandatory labels of an object may float up (confidentiality) or down (integrity) when information is written to its associated atomic object.

This is specified using the "floating\_confidentiality\_level" and "floating\_integrity\_level" attributes of the executing process, which have four possible values:

- NO\_FLOAT: switches off the floating mechanism;
- FLOAT\_IN: enables the process's mandatory context to float;
- FLOAT\_OUT: enables the object's mandatory labels to float;
- FLOAT\_IN\_OUT: enables both to float.

If the floating of the mandatory context of a process *P* is enabled (FLOAT\_IN and FLOAT\_IN\_OUT), then when information is read from the atomic object associated with an object *A*:

- if LABEL\_DOMINATES (confidentiality context of *P*, confidentiality label of *A*) is **false** then the new confidentiality context is given by FLOAT\_UPGRADE (confidentiality context of *P*, confidentiality label of *A*);
- if LABEL\_DOMINATES (integrity label of *A*, integrity context of *P*) is **false**, then the new integrity context is given by FLOAT\_DOWNGRADE (integrity context of *P*, integrity label of *A*).

If the floating of an object's mandatory labels is enabled (FLOAT\_OUT and FLOAT\_IN\_OUT), then when the atomic object associated with an object *A* is written to:

- if LABEL\_DOMINATES (confidentiality label of *A*, confidentiality context of the calling process) is **false** then the new confidentiality label is given by FLOAT\_UPGRADE (confidentiality label of *A*, confidentiality context of the calling process);
- if LABEL\_DOMINATES (integrity context of the calling process, integrity label of *A*) is **false** then the new integrity label is given by FLOAT\_DOWNGRADE (integrity label of *A*, integrity context of the calling process).

FLOAT\_UPGRADE and FLOAT\_DOWNGRADE are defined as follows:

```

FLOAT_UPGRADE (
    upgradable_label : Security_label,
    higher_label     : Security_label
)
    upgraded_label : Security_label
  
```

*upgraded\_label* is the conjunction of *upgradable\_label* and *higher\_label* unless *upgradable\_label* is null in which case *upgraded\_label* is *higher\_label*.

```

FLOAT_DOWNGRADE (
    downgradable_label : Security_label,
    lower_label        : Security_label
)
    downgraded_label : Security_label
  
```

*downgraded\_label* is the disjunction of *downgradable\_label* and *lower\_label* unless *lower\_label* is null in which case *downgraded\_label* is also null.

The floating of mandatory labels requires the process to have the CONTROL\_MANDATORY right to the object.

The confidentiality context of a process *process* is subject to the constraints:

- It must be dominated by the user confidentiality clearance.

- It must lie within the confidentiality range of the workstation i.e. CONFIDENTIALITY\_LABEL\_WITHIN\_RANGE (*process, station*) must be **true**.

The confidentiality label of an object A must lie within the confidentiality range of the volume V in which it resides, i.e. CONFIDENTIALITY\_LABEL\_WITHIN\_RANGE (A, V) must be **true**.

The integrity context must continue to lie within the integrity range of the workstation on which the process is running i.e. INTEGRITY\_LABEL\_WITHIN\_RANGE (*process, station*) must be **true**.

The integrity label of an object A must lie within the integrity range of the volume V in which it resides, i.e. INTEGRITY\_LABEL\_WITHIN\_RANGE (A, V) must be **true**.

#### NOTES

- 1 If any of the above conditions results in the process's mandatory context or the object's mandatory label not being changed, then reading and writing of the object are forbidden, as defined in 20.1.4.
- 2 CONTROL\_MANDATORY right is required for label changes to be effected either explicitly using OBJECT\_SET\_CONFIDENTIALITY\_LABEL and OBJECT\_SET\_INTEGRITY\_LABEL or implicitly using floating security labels.
- 3 In order to determine whether these constraints have been violated, access must be made to the objects involved i.e. the user, the station and the volume. These accesses are not also subject to mandatory access control, which could lead to the further floating of the mandatory context of the current process. These accesses constitute additional bitwise read accesses which are intrinsic covert channels to PCTE (see 20.1.8.2) and are permitted.
- 4 An object of type "process" (or a descendant type) cannot have its mandatory labels changed by output floating, regardless of the process status. An operation which tries to write to such an object and would cause floating fails with the relevant confinement violation error.

### 20.1.7 Implementation restrictions

A trusted implementation of PCTE may have implementation-defined restrictions on various aspects of the security model. In particular there may be implementation-defined restrictions of the following kinds:

- restrictions on the number of confidentiality classes (0 or more);
- restrictions on the number of integrity classes (0 or more);
- restrictions on the form of the confidentiality labels, e.g. may not allow a disjunction;
- restrictions on the form of the integrity labels, e.g. may not allow a conjunction;
- restrictions on creation of links between levels (e.g. may not allow any links to cross differently labelled objects for designated information classes).

In some implementations there may be predefined classes. These predefined classes may be protected using particular implementation-defined techniques.

### 20.1.8 Built-in policy aspects

Some aspects of the security policy of any PCTE environment are enforced by the PCTE interfaces. Any attempt to violate the built-in policy aspect raises the error condition SECURITY\_POLICY\_WOULD\_BE\_VIOLATED.



### 20.1.8.1 Protection of predefined SDSs

The predefined SDSs "system", "discretionary\_security", "mandatory\_security", "metasds" and "accounting" have to be protected against any modification.

Thus, for all these SDSs, the atomic and composite ACLs contain an entry corresponding to the predefined security group ALL\_USERS - which is automatically set in the discretionary context of all processes - with WRITE\_ATTRIBUTES, WRITE\_LINKS, APPEND\_LINKS and DELETE access DENIED. The other access rights are set to UNDEFINED.

NOTE – Any attempt by clause 10 operations to change a predefined SDS is forbidden.

### 20.1.8.2 Covert channels

A *covert channel* is a communication channel that allows a process to transfer information in a manner which violates the system's security policy. The mandatory and discretionary security conditions defined in previous clauses are enforced throughout PCTE. An appropriate error condition is raised whenever a given operation would result in a violation of such rules and of the other aspects of the built-in policy.

Two kinds of access are identified for the purposes of mandatory security:

- *data access*: accesses of this kind are implied when data items are explicitly transferred between a process and an object.
- *bitwise access*: accesses of this kind are implied when the status of an object (or of a process) is modified or queried as a side effect of an operation. The term "status" is used here as opposed to the data values held in the object and which can be manipulated via the data accesses defined above.

A bitwise read access is:

- an integrity covert channel where the process strictly dominates the object in integrity;
- a confidentiality covert channel where the object strictly dominates the process in confidentiality.

A bitwise write access is:

- a confidentiality covert channel where the process strictly dominates the object in confidentiality;
- an integrity covert channel where the object strictly dominates the process in integrity.

Both kinds of access imply transfer of information between processes and objects (or other processes). However, in the built-in policy, control of information flow is dealt with differently for the two kinds of access:

- all "data accesses" must conform to the mandatory security rules as defined earlier in this major clause;
- a certain number of "bitwise accesses" are allowed which would otherwise violate the security rules. These are classified as intrinsic covert channels. PCTE implementations can restrict information flow through covert channels. The events leading to intrinsic covert channels are all those associated with bitwise write accesses.

The following operations imply bitwise read access:

- LOCK\_SET\_OBJECT, LOCK\_UNSET\_OBJECT, LOCK\_SET\_INTERNAL\_MODE, LOCK\_RESET\_INTERNAL\_MODE;
- any access to an object which implies a check on access rights.

The following operations imply bitwise write access:

- LOCK\_SET\_OBJECT, LOCK\_UNSET\_OBJECT, LOCK\_SET\_INTERNAL\_MODE, LOCK\_RESET\_INTERNAL\_MODE;
- ACTIVITY\_START, ACTIVITY\_ABORT, ACTIVITY\_END;
- MESSAGE\_RECEIVE\_NO\_WAIT, MESSAGE\_RECEIVE\_WAIT, MESSAGE\_PEEK if the message queue is full;
- LINK\_CREATE (creation of an implicit link);
- any write to the audit file;
- any write to the accounting log;
- any implicit creation or deletion of a usage designation link;
- any operation which creates or deletes an object (creation or deletion of an "object\_on\_volume" link);
- OBJECT\_MOVE, on the destinations of external non-designation links of the object (if moved) and each moved component.

## 20.2 Operations for mandatory security operation

### 20.2.1 DEVICE\_SET\_CONFIDENTIALITY\_RANGE

```

DEVICE_SET_CONFIDENTIALITY_RANGE (
    device      : Device_designator,
    high_label  : Security_label,
    low_label   : Security_label
)

```

DEVICE\_SET\_CONFIDENTIALITY\_RANGE sets the confidentiality range high label and confidentiality range low label of *device* to *high\_label* and *low\_label* respectively, subject to the following conditions, where *device'* is *device* with the confidentiality range so changed, *station* is the workstation controlling *device*, *simply\_enlarged* is CONFIDENTIALITY\_RANGE\_WITHIN\_RANGE (*device*, *device'*), and *simply\_reduced* is CONFIDENTIALITY\_RANGE\_WITHIN\_RANGE (*device'*, *device*):

- If *simply\_enlarged* or not *simply\_reduced*, then CONFIDENTIALITY\_RANGE\_WITHIN\_RANGE (*device'*, *station*) must be **true**.
- If *simply\_reduced* or not *simply\_enlarged*, and there is a volume *volume* mounted on *device*, then CONFIDENTIALITY\_RANGE\_WITHIN\_RANGE (*volume*, *device'*) must be **true**.

If floating of security labels is switched on for the calling process, the facility is ignored for this operation.

A write lock of the default mode is obtained on *device*.

**Errors**

ACCESS\_ERRORS (*device*, ATOMIC, CHANGE, CONTROL\_MANDATORY)

DEVICE\_IS\_UNKNOWN (*device*)

CONFIDENTIALITY\_LABEL\_IS\_INVALID (*high\_label*)

CONFIDENTIALITY\_LABEL\_IS\_INVALID (*low\_label*)

LABEL\_RANGE\_IS\_BAD (*high\_label*, *low\_label*)

PROCESS\_IS\_IN\_TRANSACTION

RANGE\_IS\_OUTSIDE\_RANGE (*object*, *station*)

If there is a volume *volume* mounted on the device:

RANGE\_IS\_OUTSIDE\_RANGE (*volume*, *object*)

NOTE - It is possible that the range is being enlarged and reduced at the same time, e.g. if both *high\_label* and *low\_label* are upgrades, in which case all relevant constraints must be applied.

**20.2.2 DEVICE\_SET\_INTEGRITY\_RANGE**

```

DEVICE_SET_INTEGRITY_RANGE (
    device      : Device_designator,
    high_label  : Security_label,
    low_label   : Security_label
)

```

DEVICE\_SET\_INTEGRITY\_RANGE sets the integrity range high label and integrity range low label of *device* to *high\_label* and *low\_label* respectively, subject to the following conditions, where *device'* is *device* with the integrity range so changed, *station* is the workstation controlling *device*, *simply\_enlarged* is INTEGRITY\_RANGE\_WITHIN\_RANGE (*device*, *device'*) and *simply\_reduced* is INTEGRITY\_RANGE\_WITHIN\_RANGE (*device'*, *device*):

- If *simply\_enlarged* or not *simply\_reduced*, then INTEGRITY\_RANGE\_WITHIN\_RANGE (*device'*, *station*) must be **true**.
- If *simply\_reduced* or not *simply\_enlarged*, and there is a volume *volume* mounted on *device*, then INTEGRITY\_RANGE\_WITHIN\_RANGE (*volume*, *device'*) must be **true**.

If floating of security labels is switched on for the calling process, the facility is ignored for this operation.

A write lock of the default mode is obtained on *device*.

**Errors**

ACCESS\_ERRORS (*device*, ATOMIC, CHANGE, CONTROL\_MANDATORY)

DEVICE\_IS\_UNKNOWN (*device*)

INTEGRITY\_LABEL\_IS\_INVALID (*high\_label*)

INTEGRITY\_LABEL\_IS\_INVALID (*low\_label*)

LABEL\_RANGE\_IS\_BAD (*high\_label*, *low\_label*)

PROCESS\_IS\_IN\_TRANSACTION

RANGE\_IS\_OUTSIDE\_RANGE (*object*, *station*)

If there is a volume *volume* mounted on the device:

RANGE\_IS\_OUTSIDE\_RANGE (*volume*, *object*)

NOTE - It is possible that the range is being enlarged and reduced at the same time, e.g. if both *high\_label* and *low\_label* are upgrades, in which case all relevant constraints must be applied.

### 20.2.3 EXECUTION\_SITE\_SET\_CONFIDENTIALITY\_RANGE

```
EXECUTION_SITE_SET_CONFIDENTIALITY_RANGE (
    execution_site : Execution_site_designator,
    high_label     : Security_label,
    low_label      : Security_label
)
```

EXECUTION\_SITE\_SET\_CONFIDENTIALITY\_RANGE sets the confidentiality range high label and confidentiality range low label of *execution\_site* to *high\_label* and *low\_label* respectively, subject to the following conditions, where *execution\_site* is *execution\_site* with the confidentiality range so changed.

If CONFIDENTIALITY\_RANGE\_WITHIN\_RANGE (*execution\_site*', *execution\_site*) or not CONFIDENTIALITY\_RANGE\_WITHIN\_RANGE (*execution\_site*, *execution\_site*):

- for each device D controlled by *execution\_site*,  
CONFIDENTIALITY\_RANGE\_WITHIN\_RANGE (D, *execution\_site*') is **true**.
- for each process P executing on *execution\_site*,  
CONFIDENTIALITY\_LABEL\_WITHIN\_RANGE (P, *execution\_site*') is **true**.

If floating of security labels is switched on for the calling process, the facility is ignored for this operation.

A write lock of the default mode is established on *execution\_site*.

#### Errors

ACCESS\_ERRORS (*execution\_site*, ATOMIC, CHANGE, CONTROL\_MANDATORY)

DEVICE\_IS\_UNKNOWN (*execution\_site*)

CONFIDENTIALITY\_LABEL\_IS\_INVALID (*high\_label*)

CONFIDENTIALITY\_LABEL\_IS\_INVALID (*low\_label*)

LABEL\_IS\_OUTSIDE\_RANGE (D, *execution\_site*)

LABEL\_IS\_OUTSIDE\_RANGE (P, *execution\_site*)

RANGE\_IS\_BAD (*high\_label*, *low\_label*)

PROCESS\_IS\_IN\_TRANSACTION

### 20.2.4 EXECUTION\_SITE\_SET\_INTEGRITY\_RANGE

```
EXECUTION_SITE_SET_INTEGRITY_RANGE (
    execution_site : Execution_site_designator,
    high_label     : Security_label,
    low_label      : Security_label
)
```

EXECUTION\_SITE\_SET\_INTEGRITY\_RANGE sets the integrity range high label and integrity range low label of *execution\_site* to *high\_label* and *low\_label* respectively, subject to the following conditions, where *execution\_site'* is *execution\_site* with the integrity range so changed.

If INTEGRITY\_RANGE\_WITHIN\_RANGE (*execution\_site'*, *execution\_site*) or not INTEGRITY\_RANGE\_WITHIN\_RANGE (*execution\_site*, *execution\_site'*):

- for each device D controlled by *execution\_site*, INTEGRITY\_RANGE\_WITHIN\_RANGE (D, *execution\_site'*) is **true**.
- for each process P executing on *execution\_site*, INTEGRITY\_LABEL\_WITHIN\_RANGE (P, *execution\_site'*) is **true**.

If floating of security labels is switched on for the calling process, the facility is ignored for this operation.

A write lock of the default mode is established on *execution\_site*.

### Errors

ACCESS\_ERRORS (*execution\_site*, ATOMIC, CHANGE, CONTROL\_MANDATORY)

DEVICE\_IS\_UNKNOWN (*execution\_site*)

INTEGRITY\_LABEL\_IS\_INVALID (*high\_label*)

INTEGRITY\_LABEL\_IS\_INVALID (*low\_label*)

LABEL\_IS\_OUTSIDE\_RANGE (D, *execution\_site*)

LABEL\_IS\_OUTSIDE\_RANGE (P, *execution\_site*)

LABEL\_RANGE\_IS\_BAD (*high\_label*, *low\_label*)

PROCESS\_IS\_IN\_TRANSACTION

## 20.2.5 OBJECT\_SET\_CONFIDENTIALITY\_LABEL

```
OBJECT_SET_CONFIDENTIALITY_LABEL (
    object : Object_designator,
    label  : Security_label
)
```

OBJECT\_SET\_CONFIDENTIALITY\_LABEL sets the confidentiality label of *object* to *label*.

If the previous value of the confidentiality label of *object* is L, then RELATIVE\_LABEL\_DOMINATES\_IN\_CONFIDENTIALITY (calling process, *label*, L) must be **true**.

CONFIDENTIALITY\_LABEL\_WITHIN\_RANGE (*object*, *volume*) must remain **true**, where *volume* is the volume on which the *object* resides.

If floating of security labels is switched on for the calling process, the facility is ignored for this operation.

A write lock of the default mode is obtained on the designated *object*.

### Errors

ACCESS\_ERRORS (*object*, ATOMIC, CHANGE, CONTROL\_MANDATORY)

CONFIDENTIALITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED (*object*, ATOMIC)  
 CONFIDENTIALITY\_LABEL\_IS\_INVALID (*label*)  
 LABEL\_IS\_OUTSIDE\_RANGE (*object*, *volume*)  
 OBJECT\_IS\_A\_PROCESS (*object*)  
 OBJECT\_LABEL\_CANNOT\_BE\_CHANGED\_IN\_TRANSACTION (*object*)

#### 20.2.6 OBJECT\_SET\_INTEGRITY\_LABEL

```
OBJECT_SET_INTEGRITY_LABEL (
    object : Object_designator,
    label   : Security_label
)
```

OBJECT\_SET\_INTEGRITY\_LABEL sets the integrity label of *object* to *label*.

If the previous value of the integrity label of *object* is L, then RELATIVE\_LABEL\_DOMINATES\_IN\_INTEGRITY (calling process, L, *label*) must be true.

INTEGRITY\_LABEL\_WITHIN\_RANGE (*object*, *volume*) must remain **true**, where *volume* is the volume on which *object* resides.

If floating of security labels is switched on for the calling process, the facility is ignored for this operation.

A write lock of the default mode is obtained on the designated *object*.

##### Errors

ACCESS\_ERRORS (*object*, ATOMIC, CHANGE, CONTROL\_MANDATORY)  
 INTEGRITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED (*object*, ATOMIC)  
 INTEGRITY\_LABEL\_IS\_INVALID (*label*)  
 LABEL\_IS\_OUTSIDE\_RANGE (*object*, *volume*)  
 OBJECT\_IS\_A\_PROCESS (*object*)  
 OBJECT\_LABEL\_CANNOT\_BE\_CHANGED\_IN\_TRANSACTION (*object*)

#### 20.2.7 VOLUME\_SET\_CONFIDENTIALITY\_RANGE

```
VOLUME_SET_CONFIDENTIALITY_RANGE (
    volume   : Volume_designator,
    high_label : Security_label,
    low_label  : Security_label
)
```

VOLUME\_SET\_CONFIDENTIALITY\_RANGE sets the confidentiality high label and confidentiality low label of *volume* to *high\_label* and *low\_label* respectively subject to the following conditions, where *volume'* is *volume* with its confidentiality range so changed, *simply\_enlarged* is CONFIDENTIALITY\_RANGE\_WITHIN\_RANGE (*volume*, *volume'*), and *simply\_reduced* is CONFIDENTIALITY\_RANGE\_WITHIN\_RANGE (*volume'*, *volume*).

- If *simply\_enlarged* or not *simply\_reduced*, let *device* be the device on which the *volume* is mounted, then CONFIDENTIALITY\_RANGE\_WITHIN\_RANGE (*volume'*, *device*) must be **true**.

- If *simply\_reduced* or not *simply\_enlarged*, then for each object G residing on the volume, CONFIDENTIALITY\_LABEL\_WITHIN\_RANGE (G, *volume'*) must be **true**.

If floating of security labels is switched on for the calling process, the facility is ignored for this operation.

A write lock of the default mode is obtained on *volume*.

### Errors

ACCESS\_ERRORS (*volume*, ATOMIC, CHANGE, CONTROL\_MANDATORY)

CONFIDENTIALITY\_LABEL\_IS\_INVALID (*high\_label*)

CONFIDENTIALITY\_LABEL\_IS\_INVALID (*low\_label*)

For each object G residing on *volume*:

LABEL\_IS\_OUTSIDE\_RANGE (G, *device*)

LABEL\_RANGE\_IS\_BAD (*high\_label*, *low\_label*)

RANGE\_IS\_OUTSIDE\_RANGE (*volume*, *device*)

PROCESS\_IS\_IN\_TRANSACTION

VOLUME\_HAS\_OBJECT\_OUTSIDE\_RANGE (*volume*, *high\_label*, *low\_label*)

VOLUME\_IS\_UNKNOWN (*volume*)

NOTE - It is possible that the range is being enlarged and reduced at the same time, e.g. if both *high\_label* and *low\_label* are upgrades, in which case both constraints must be applied.

## 20.2.8 VOLUME\_SET\_INTEGRITY\_RANGE

```
VOLUME_SET_INTEGRITY_RANGE (
    volume      : Volume_designator,
    high_label  : Security_label,
    low_label   : Security_label
)
```

VOLUME\_SET\_INTEGRITY\_RANGE sets the integrity range high label and integrity range low label of *volume* to *high\_label* and *low\_label* respectively subject to the following conditions, where *volume'* is *volume* with its integrity range so changed, *simply\_enlarged* is INTEGRITY\_RANGE\_WITHIN\_RANGE (*volume*, *volume'*), and *simply\_reduced* is INTEGRITY\_RANGE\_WITHIN\_RANGE (*volume'*, *volume*):

- If *simply\_enlarged* or not *simply\_reduced*, let *device* be the device on which the *volume* is mounted, then INTEGRITY\_RANGE\_WITHIN\_RANGE (*volume'*, *device*) must be **true**.
- If *simply\_reduced* or not *simply\_enlarged*, then for each object G residing on the volume, INTEGRITY\_LABEL\_WITHIN\_RANGE (G, *volume'*) must be **true**.

If floating of security labels is switched on for the calling process, the facility is ignored for this operation.

A write lock of the default mode is obtained on *volume*.

### Errors

ACCESS\_ERRORS (*volume*, ATOMIC, CHANGE, CONTROL\_MANDATORY)

INTEGRITY\_LABEL\_IS\_INVALID (*high\_label*)



INTEGRITY\_LABEL\_IS\_INVALID (*low\_label*)

For each object *G* residing on *volume*:

LABEL\_IS\_OUTSIDE\_RANGE (*G*, *device*)

LABEL\_RANGE\_IS\_BAD (*high\_label*, *low\_label*)

RANGE\_IS\_OUTSIDE\_RANGE (*volume*, *device*)

PROCESS\_IS\_IN\_TRANSACTION

VOLUME\_HAS\_OBJECT\_OUTSIDE\_RANGE (*volume*, *high\_label*, *low\_label*)

VOLUME\_IS\_UNKNOWN (*volume*)

NOTE - It is possible that the range is being enlarged and reduced at the same time, e.g. if both *high\_label* and *low\_label* are upgrades, in which case both constraints must be applied.

## 20.3 Mandatory security administration operations

### 20.3.1 CONFIDENTIALITY\_CLASS\_INITIALIZE

```
CONFIDENTIALITY_CLASS_INITIALIZE (
    object          : Confidentiality_class_designator,
    class_name      : Name,
    to_be_dominated : [ Confidentiality_class_designator ]
)
```

CONFIDENTIALITY\_CLASS\_INITIALIZE initializes *object* as a confidentiality class. A "known\_mandatory\_class" link keyed by *class\_name* is created from the master of the mandatory directory to *object*. If *to\_be\_dominated* is supplied, a "dominates\_in\_confidentiality" link is created from *object* to *to\_be\_dominated*, and a "confidentiality\_dominator" link is created from *to\_be\_dominated* to *object*.

If *to\_be\_dominated* is not supplied, the operation creates a new confidentiality tower consisting of the one confidentiality class *object*. If *to\_be\_dominated* is supplied, the operation adds *object* to the tail (the 'top') of an existing confidentiality tower.

Write locks of the default mode are obtained on the created links.

#### Errors

ACCESS\_ERRORS (the mandatory directory, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*object*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

If *to\_be\_dominated* is supplied:

ACCESS\_ERRORS (*object*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*to\_be\_dominated*, ATOMIC, MODIFY, APPEND\_LINKS)

MANDATORY\_CLASS\_IS\_ALREADY\_DOMINATED (*to\_be\_dominated*)

MANDATORY\_CLASS\_IS\_UNKNOWN(*object*)

MANDATORY\_CLASS\_IS\_UNKNOWN (*to\_be\_dominated*)

MANDATORY\_CLASS\_NAME\_IS\_IN\_USE (*class\_name*)

PROCESS\_IS\_IN\_TRANSACTION

NOTE - This operation does not change any copies of the mandatory directory.

### 20.3.2 GROUP\_DISABLE\_FOR\_CONFIDENTIALITY\_DOWNGRADE

```
GROUP_DISABLE_FOR_CONFIDENTIALITY_DOWNGRADE (
    group          : User_designator | User_group_designator |
                    Program_group_designator,
    confidentiality_class : Confidentiality_class_designator
)
```

GROUP\_DISABLE\_FOR\_CONFIDENTIALITY\_DOWNGRADE deletes a "may\_downgrade" link from *group* to *confidentiality\_class* and a "downgradable\_by" link from *confidentiality\_class* to *group*.

Write locks of the default mode are obtained on the deleted links.

#### Errors

ACCESS\_ERRORS (*group*, ATOMIC, MODIFY, WRITE\_LINKS)  
 ACCESS\_ERRORS (*confidentiality\_class*, ATOMIC, MODIFY, WRITE\_LINKS)  
 MANDATORY\_CLASS\_IS\_UNKNOWN (*confidentiality\_class*)  
 SECURITY\_GROUP\_IS\_NOT\_ENABLED (*group*, *confidentiality\_class*)  
 SECURITY\_GROUP\_IS\_UNKNOWN (*group*)

### 20.3.3 GROUP\_DISABLE\_FOR\_INTEGRITY\_UPGRADE

```
GROUP_DISABLE_FOR_INTEGRITY_UPGRADE (
    group          : User_designator | User_group_designator | Program_group_designator,
    integrity_class : Confidentiality_class_designator
)
```

GROUP\_DISABLE\_FOR\_INTEGRITY\_UPGRADE deletes a "may\_upgrade" link from *group* to *integrity\_class* and an "upgradable\_by" link from *integrity\_class* to *group*.

Write locks of the default mode are obtained on the links so deleted.

#### Errors

ACCESS\_ERRORS (*group*, ATOMIC, MODIFY, WRITE\_LINKS)  
 ACCESS\_ERRORS (*integrity\_class*, ATOMIC, MODIFY, WRITE\_LINKS)  
 MANDATORY\_CLASS\_IS\_UNKNOWN (*integrity\_class*)  
 SECURITY\_GROUP\_IS\_NOT\_ENABLED (*group*, *integrity\_class*)  
 SECURITY\_GROUP\_IS\_UNKNOWN (*group*)

### 20.3.4 GROUP\_ENABLE\_FOR\_CONFIDENTIALITY\_DOWNGRADE

```
GROUP_ENABLE_FOR_CONFIDENTIALITY_DOWNGRADE (
    group          : User_designator | User_group_designator |
                    Program_group_designator,
    confidentiality_class : Confidentiality_class_designator
)
```

GROUP\_ENABLE\_FOR\_CONFIDENTIALITY\_DOWNGRADE creates a "may\_downgrade" link, keyed by the confidentiality class name of *confidentiality\_class*, from *group* to

*confidentiality\_class* and a "downgradable\_by" link, keyed by the group identifier, from *confidentiality\_class* to *group*.

Write locks of the default mode are obtained on the links so created.

#### Errors

ACCESS\_ERRORS (*group*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*confidentiality\_class*, ATOMIC, MODIFY, APPEND\_LINKS)

MANDATORY\_CLASS\_IS\_UNKNOWN (*confidentiality\_class*)

SECURITY\_GROUP\_IS\_ALREADY\_ENABLED (*group*, *confidentiality\_class*)

SECURITY\_GROUP\_IS\_UNKNOWN (*group*)

### 20.3.5 GROUP\_ENABLE\_FOR\_INTEGRITY\_UPGRADE

```
GROUP_ENABLE_FOR_INTEGRITY_UPGRADE (
    group          : User_designator | User_group_designator | Program_group_designator,
    integrity_class : Confidentiality_class_designator
)
```

GROUP\_ENABLE\_FOR\_INTEGRITY\_UPGRADE creates a "may\_upgrade" link, keyed by the integrity class name of *integrity\_class*, from *group* to *integrity\_class* and an "upgradable\_by" link, keyed by the group identifier, from *integrity\_class* to *group*.

Write locks of the default mode are obtained on the links so created.

#### Errors

ACCESS\_ERRORS (*group*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*integrity\_class*, ATOMIC, MODIFY, APPEND\_LINKS)

MANDATORY\_CLASS\_IS\_UNKNOWN (*integrity\_class*)

SECURITY\_GROUP\_IS\_ALREADY\_ENABLED (*group*, *integrity\_class*)

SECURITY\_GROUP\_IS\_UNKNOWN (*group*)

### 20.3.6 INTEGRITY\_CLASS\_INITIALIZE

```
INTEGRITY_CLASS_INITIALIZE (
    object          : Integrity_class_designator,
    class_name      : Name,
    to_be_dominated : [ Integrity_class_designator ]
)
```

INTEGRITY\_CLASS\_INITIALIZE initializes *object* as an integrity class. A "mandatory\_class" link keyed by *class\_name* is created from the master of the mandatory directory to *object*. If *to\_be\_dominated* is supplied, a "dominates\_in\_integrity" link is created from *object* to *to\_be\_dominated*, and a "integrity\_dominator" link is created from *to\_be\_dominated* to *object*.

If *to\_be\_dominated* is not supplied, the operation creates a new integrity tower consisting of the one integrity class *object*. If *to\_be\_dominated* is supplied, the operation adds *object* to the tail (the "top") of an existing integrity tower.

Write locks of the default mode are obtained on the created links.

**Errors**

ACCESS\_ERRORS (the mandatory directory, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*object*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

If *to\_be\_dominated* is supplied:

ACCESS\_ERRORS (*object*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*to\_be\_dominated*, ATOMIC, MODIFY, APPEND\_LINKS)

MANDATORY\_CLASS\_IS\_ALREADY\_DOMINATED (*to\_be\_dominated*)

MANDATORY\_CLASS\_IS\_UNKNOWN(*object*)

MANDATORY\_CLASS\_IS\_UNKNOWN (*to\_be\_dominated*)

MANDATORY\_CLASS\_NAME\_IS\_IN\_USE (*class\_name*)

PROCESS\_IS\_IN\_TRANSACTION

NOTE - This operation does not change any copies of the mandatory directory.

**20.3.7 USER\_EXTEND\_CONFIDENTIALITY\_CLEARANCE**

```

USER_EXTEND_CONFIDENTIALITY_CLEARANCE (
    user          : User_designator,
    confidentiality_class : Confidentiality_class_designator
)

```

USER\_EXTEND\_CONFIDENTIALITY\_CLEARANCE creates a "cleared\_for" link, keyed by the name of the confidentiality class *confidentiality\_class*, from *user* to *confidentiality\_class* and a "having\_clearance" link, keyed by the group identifier, from *confidentiality\_class* to *user*.

Write locks of the default mode are obtained on the links so created.

**Errors**

ACCESS\_ERRORS (*user*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*confidentiality\_class*, ATOMIC, MODIFY, APPEND\_LINKS)

MANDATORY\_CLASS\_IS\_UNKNOWN (*confidentiality\_class*)

SECURITY\_GROUP\_IS\_UNKNOWN (*user*)

USER\_IS\_ALREADY\_CLEARED\_TO\_CLASS (*user*, *confidentiality\_class*)

USER\_IS\_IN\_USE (*user*)

**20.3.8 USER\_EXTEND\_INTEGRITY\_CLEARANCE**

```

USER_EXTEND_INTEGRITY_CLEARANCE (
    user          : User_designator,
    integrity_class : Integrity_class_designator
)

```

USER\_EXTEND\_INTEGRITY\_CLEARANCE creates a "cleared\_for" link, keyed by the name of the integrity class *integrity\_class*, from *user* to *integrity\_class*, and a "having\_clearance" link, keyed by the group identifier, from *integrity\_class* to *user*.

Write locks of the default mode are obtained on the links so created.

**Errors**

ACCESS\_ERRORS (*user*, ATOMIC, MODIFY, APPEND\_LINKS)  
 ACCESS\_ERRORS (*integrity\_class*, ATOMIC, MODIFY, APPEND\_LINKS)  
 MANDATORY\_CLASS\_IS\_UNKNOWN (*integrity\_class*)  
 SECURITY\_GROUP\_IS\_UNKNOWN (*user*)  
 USER\_IS\_ALREADY\_CLEARED\_TO\_CLASS (*user*, *integrity\_class*)  
 USER\_IS\_IN\_USE (*user*)

**20.3.9 USER\_REDUCE\_CONFIDENTIALITY\_CLEARANCE**

```

USER_REDUCE_CONFIDENTIALITY_CLEARANCE (
    user                : User_designator,
    confidentiality_class : Confidentiality_class_designator
)

```

USER\_REDUCE\_CONFIDENTIALITY\_CLEARANCE deletes a "cleared\_for" link from *user* to *confidentiality\_class* or to a confidentiality class which dominates *confidentiality\_class* and a "having\_clearance" link from that confidentiality class to *user*.

Write locks of the default mode are obtained on the links so deleted.

**Errors**

ACCESS\_ERRORS (*user*, ATOMIC, MODIFY, WRITE\_LINKS)  
 ACCESS\_ERRORS (*confidentiality\_class*, ATOMIC, MODIFY, WRITE\_LINKS)  
 MANDATORY\_CLASS\_IS\_UNKNOWN (*confidentiality\_class*)  
 SECURITY\_GROUP\_IS\_UNKNOWN (*user*)  
 USER\_IS\_NOT\_CLEARED\_TO\_CLASS (*user*, *confidentiality\_class*)  
 USER\_IS\_IN\_USE (*user*)

NOTE - There is at most one link that satisfies the conditions above for deletion.

**20.3.10 USER\_REDUCE\_INTEGRITY\_CLEARANCE**

```

USER_REDUCE_INTEGRITY_CLEARANCE (
    user                : User_designator,
    integrity_class     : Integrity_class_designator
)

```

USER\_REDUCE\_INTEGRITY\_CLEARANCE deletes a "cleared\_for" link from *user* to *integrity\_class* or to an integrity class which dominates *integrity\_class* and a "having\_clearance" link from that integrity class to *user*.

Write locks of the default mode are obtained on the deleted links.

**Errors**

ACCESS\_ERRORS (*user*, ATOMIC, MODIFY, WRITE\_LINKS)  
 ACCESS\_ERRORS (*integrity\_class*, ATOMIC, MODIFY, WRITE\_LINKS)  
 MANDATORY\_CLASS\_IS\_UNKNOWN (*integrity\_class*)  
 SECURITY\_GROUP\_IS\_UNKNOWN (*user*)

USER\_IS\_NOT\_CLEARED\_TO\_CLASS (*user*, *integrity\_class*)

USER\_IS\_IN\_USE (*user*)

## 20.4 Mandatory security operations for processes

### 20.4.1 PROCESS\_SET\_CONFIDENTIALITY\_LABEL

```
PROCESS_SET_CONFIDENTIALITY_LABEL (
    process           : [ Process_designator ],
    confidentiality_label : Security_label
)
```

If no value is supplied for *process*, *process* designates the calling process.

PROCESS\_SET\_CONFIDENTIALITY\_LABEL sets the confidentiality label of *process* to *confidentiality\_label*.

If floating of security labels is switched on for the calling process, the facility is ignored for this operation.

#### Errors

If *process* is not the calling process:

ACCESS\_ERRORS (*process*, ATOMIC, CHANGE, CONTROL\_MANDATORY)

CONFIDENTIALITY\_LABEL\_IS\_INVALID (*confidentiality\_label*)

If *process* is the calling process:

LABEL\_IS\_OUTSIDE\_RANGE (*process*, execution site of *process*)

LABEL\_IS\_OUTSIDE\_RANGE (*process*, volume on which *process* resides)

PROCESS\_CONFIDENTIALITY\_IS\_NOT\_DOMINATED (*confidentiality\_label*, *process*)

If *process* is not the calling process:

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, READY)

PROCESS\_IS\_UNKNOWN (*process*)

USER\_IS\_NOT\_CLEARED (*process*, *confidentiality\_label*)

### 20.4.2 PROCESS\_SET\_FLOATING\_CONFIDENTIALITY\_LEVEL

```
PROCESS_SET_FLOATING_CONFIDENTIALITY_LEVEL (
    process           : [ Process_designator ],
    floating_mode      : Floating_level
)
```

If no value is supplied for *process*, *process* designates the calling process.

PROCESS\_SET\_FLOATING\_CONFIDENTIALITY\_LEVEL sets the floating confidentiality level of *process* to *floating\_mode*.

#### Errors

If *process* is not the calling process:

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, WRITE\_ATTRIBUTES)

If *process* is not the calling process:

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, READY)

PROCESS\_IS\_UNKNOWN (*process*)

### 20.4.3 PROCESS\_SET\_FLOATING\_INTEGRITY\_LEVEL

```
PROCESS_SET_FLOATING_INTEGRITY_LEVEL (
    process          : [ Process_designator ],
    floating_mode    : Floating_level
)
```

If no value is supplied for *process*, *process* designates the calling process.

PROCESS\_SET\_FLOATING\_INTEGRITY\_LEVEL sets the floating integrity level of *process* to *floating\_mode*.

#### Errors

If *process* is not the calling process:

ACCESS\_ERRORS (*process*, ATOMIC, MODIFY, WRITE\_ATTRIBUTES)

If *process* is not the calling process:

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, READY)

PROCESS\_IS\_UNKNOWN (*process*)

### 20.4.4 PROCESS\_SET\_INTEGRITY\_LABEL

```
PROCESS_SET_INTEGRITY_LABEL (
    process          : [ Process_designator ],
    integrity_label : Security_label
)
```

If no value is supplied for *process*, *process* designates the calling process.

PROCESS\_SET\_INTEGRITY\_LABEL sets the integrity label of *process* to *integrity\_label*.

If floating of security labels is switched on for the calling process, the facility is ignored for this operation.

#### Errors

If *process* is not the calling process:

ACCESS\_ERRORS (*process*, ATOMIC, CHANGE, CONTROL\_MANDATORY)

PROCESS\_INTEGRITY\_DOES\_NOT\_DOMINATE (*integrity\_label*, *process*)

INTEGRITY\_LABEL\_IS\_INVALID (*integrity\_label*)

If *process* is the calling process:

LABEL\_IS\_OUTSIDE\_RANGE (*process*, execution site of *process*)

LABEL\_IS\_OUTSIDE\_RANGE (*process*, volume on which *process* resides)

If *process* is not the calling process:

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, READY)

PROCESS\_IS\_UNKNOWN (*process*)



## 21 Auditing

### 21.1 Auditing concepts

#### 21.1.1 Audit files

```

Selectable_event_type = WRITE | READ | COPY | ACCESS_CONTENTS | EXPLOIT
| CHANGE_ACCESS_CONTROL_LIST | CHANGE_LABEL
| USE_PREDEFINED_GROUP | SET_USER_IDENTITY
| WRITE_CONFIDENTIALITY_VIOLATION | READ_CONFIDENTIALITY_VIOLATION
| WRITE_INTEGRITY_VIOLATION | READ_INTEGRITY_VIOLATION | COVERT_CHANNEL
| INFORMATION

Mandatory_event_type = CHANGE_IDENTIFICATION | SELECT_AUDIT_EVENT
| SECURITY_ADMINISTRATION

Auditing_record = Object_auditing_record
| Exploit_auditing_record
| Information_auditing_record
| Copy_auditing_record
| Security_auditing_record

Basic_auditing_record ::
    USER          : Group_identifier
    TIME           : Time
    WORKSTATION    : Execution_site_identifier
    EVENT_TYPE     : Selectable_event_type | Mandatory_event_type
    RETURN_CODE    : Return_code
    PROCESS        : Exact_identifier

Object_auditing_record :: Basic_auditing_record &&
    OBJECT        : Exact_identifier

Exploit_auditing_record :: Basic_auditing_record &&
    NEW_PROCESS   : Exact_identifier
    EXPLOITED_OBJECT : Exact_identifier

Information_auditing_record :: Basic_auditing_record &&
    INFORMATION   : String

Copy_auditing_record :: Basic_auditing_record &&
    SOURCE        : Exact_identifier
    DESTINATION   : Exact_identifier

Security_auditing_record :: Basic_auditing_record &&
    GROUP         : Exact_identifier

Exact_identifier = Text

Audit_file = seq of Auditing_record

Return_code = FAILURE | SUCCESS

sds discretionary_security:

import object type system-object, system-workstation;

audit_file: child type of object with
contents audit_file;
link
    audit_of: reference link (number) to workstation reverse audit;
end audit_file;

```

**extend object type workstation with link**

audit: (**navigate**) **existence link to** audit\_file **reverse** audit\_of;  
**end** workstation;

**end** discretionary\_security;

An audit file is an object which stores data associated with events that occur on one or more workstations. It may be associated with one or more workstations which share the same administration volume. The audit file associated with a workstation is the destination of an "audit" link from the workstation.

The audit file contains auditing records, each of which records information concerning one event on the workstation. An auditing record has a general part and a part that depends on the event type of the event being audited.

The general part, represented by the fields of the basic auditing record, is defined as follows:

- USER: the identity of the user invoking the operation giving rise to the event;
- TIME: the system time of the event;
- WORKSTATION: the workstation on which the event takes place;
- EVENT\_TYPE: the event type of the event;
- RETURN\_CODE: FAILURE if the operation giving rise to the event terminates in an error, SUCCESS otherwise;
- PROCESS: the process performing the operation giving rise to the event.

Event-type-specific fields are defined as follows:

- Events of type SELECT\_AUDIT\_EVENT are represented by basic auditing records;
- For object auditing records, representing events of types WRITE, READ, ACCESS\_CONTENTS, CHANGE\_ACCESS\_CONTROL\_LIST, CHANGE\_LABEL, WRITE\_CONFIDENTIALITY\_VIOLATION, WRITE\_INTEGRITY\_VIOLATION, READ\_CONFIDENTIALITY\_VIOLATION, READ\_INTEGRITY\_VIOLATION, SECURITY\_ADMINISTRATION, and COVERT\_CHANNEL:
  - . OBJECT: the object on which the operation takes place.
- For exploit auditing records, representing events of type EXPLOIT:
  - . NEW\_PROCESS: the process resulting from the exploitation of the object, e.g. if the operation has started execution of a program;
  - . EXPLOITED\_OBJECT: the object being exploited.
- For information auditing records, representing events of type INFORMATION:
  - . INFORMATION: the message associated with the event.
- For copy auditing records, representing events of types COPY and CHANGE\_IDENTIFICATION:
  - . SOURCE: the object being copied from, or the old identification of the object;
  - . DESTINATION: the object being copied to, or the new identification of the object.

- For security auditing records, representing events of types USE\_PREDEFINED\_GROUP, and SET\_USER\_IDENTITY:
  - . GROUP: the group being used, the user identifier being set or the user performing the audit selection.

If, when writing to the audit file, the write fails because the audit file is unavailable for some reason, then the operation which caused the auditable event to occur waits until an audit file is made available, unless the calling process is acting with the predefined group PCTE\_AUDIT. The means by which the audit file unavailability is notified to the operators of the PCTE installation is implementation-defined.

#### NOTES

- 1 The usage mode of the "audit" link type prevents any create or delete accesses. It is the role of an implementation-dependent bootstrap procedure to ensure that the audit file exists on a workstation when it is brought up. The audit data must be protected so that access to it is limited to users who are authorized for audit data.
- 2 No constraints on the label of the audit file are enforced by the system when the system writes to the audit file (i.e. it is up to the auditor to define it). When the system writes to the audit file, a bitwise write occurs but even in the case where this bitwise write results in a covert channel, it is not audited.

### 21.1.2 Audit selection criteria

General\_criterion = Selectable\_event\_type \* Selected\_return\_code

User\_criterion = Selectable\_event\_type \* Group\_identifier

Confidentiality\_criterion = Selectable\_event\_type \* Security\_label

Integrity\_criterion = Selectable\_event\_type \* Security\_label

Object\_criterion = Selectable\_event\_type \* Object\_designator

Audit\_status = ENABLED | DISABLED

Selection\_criterion = General\_criterion | Specific\_criterion

Specific\_criterion = User\_criterion | Confidentiality\_criterion | Integrity\_criterion | Object\_criterion

Removed\_criterion = Selectable\_event\_type | Specific\_criterion

Selected\_return\_code = Return\_code | ANY\_CODE

Criterion\_type = GENERAL | USER\_DEPENDENT | CONFIDENTIALITY\_DEPENDENT | INTEGRITY\_DEPENDENT | OBJECT\_DEPENDENT

General\_criteria = **set of** General\_criterion

User\_criteria = **set of** User\_criterion

Confidentiality\_criteria = **set of** Confidentiality\_criterion

Integrity\_criteria = **set of** Integrity\_criterion

Object\_criteria = **set of** Object\_criterion

Criteria = General\_criteria | User\_criteria | Confidentiality\_criteria | Integrity\_criteria | Object\_criteria

Event types may be *selected* for auditing on a per workstation basis. When a selected event occurs, audit data is written to the audit file associated with the workstation where the event occurred. The event types CHANGE\_IDENTIFICATION, SELECT\_AUDIT\_EVENT and

SECURITY\_ADMINISTRATION are always audited, regardless of the current selection criteria. A list of event types is in annex E.

Selected events are only audited when auditing is *enabled* on the workstation. When auditing is *disabled*, only the event types that are always audited are audited.

Events are selected on the basis of their types and either a return code, a user, an object, or a label. Each workstation maintains a set of *audit selection criteria*. The set of audit selection criteria is not persistent across workstation failure.

Criteria of each type select events as follows:

- General criterion: all events of the specified type and with the specified return code are selected for auditing, or if the specified return code is ANY\_CODE then all events of that type are selected.
- User-dependent criterion: all events of the specified type and being performed on behalf of the user identified by the group identifier are selected for auditing.
- Confidentiality-dependent criterion: all events of the specified type that are performed on objects of the specified confidentiality label are selected for auditing.
- Integrity-dependent criterion: all events of the specified type that are performed on objects of the specified integrity label are selected for auditing.
- Object-dependent criterion: all events of the specified type that are performed on the specified object are selected for auditing.

## 21.2 Auditing operations

### 21.2.1 AUDIT\_ADD\_CRITERION

```
AUDIT_ADD_CRITERION (
    station      : Workstation_designator,
    criterion    : Selection_criterion
)
```

AUDIT\_ADD\_CRITERION adds the criterion *criterion* to the audit selection criteria for the workstation *station*. Events of the type specified in *criterion* will then be audited on *station*, dependent on the type of *criterion* specified:

- General criterion: The events are recorded on the basis of the return code of the operation generating the event. If the event type is already selected with the same return code, then the operation has no effect.
- Confidentiality-dependent criterion: Events performed on objects of the specified confidentiality label are audited on *station*. If the event type and confidentiality label are already selected then the operation has no effect.
- Integrity-dependent criterion: Events performed on objects of the specified integrity label are be audited on *station*. If the event type and integrity label are already selected then the operation has no effect.
- Object-dependent criterion: Events performed on the specified object are audited on *station*. The object specified by *criterion* must be accessible. If the event type and object are already selected then the operation has no effect.

- User-dependent criterion: Events performed by the specified user are audited on *station*. If the event type and user are already selected then the operation has no effect.

### Errors

For confidentiality-dependent criterion:

CONFIDENTIALITY\_LABEL\_IS\_INVALID (security label specified by *criterion*)

For object-dependent criterion:

DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED (specified object, ATOMIC)

For user-dependent criterion:

GROUP\_IDENTIFIER\_IS\_INVALID (group identifier of *criterion*)

USER\_IS\_UNKNOWN (user specified by *criterion*)

For integrity-dependent criterion:

INTEGRITY\_LABEL\_IS\_INVALID (security label specified by *criterion*)

OBJECT\_IS\_INACCESSIBLE (*station*, ATOMIC)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_AUDIT)

WORKSTATION\_IS\_UNKNOWN (*station*)

## 21.2.2 AUDIT\_FILE\_COPY\_AND\_RESET

```
AUDIT_FILE_COPY_AND_RESET (
    source      : Audit_file_designator,
    destination : Audit_file_designator
)
```

AUDIT\_FILE\_COPY\_AND\_RESET copies the audit file *source* into the audit file *destination*. The contents of *source* is cleared. No audit records are lost.

This operation may not be invoked from within a transaction.

Write locks of the default mode are obtained on *source*, on *destination*, and on the created and deleted links.

### Errors

ACCESS\_ERRORS (*source*, ATOMIC, MODIFY, (READ\_CONTENTS, WRITE\_CONTENTS))

ACCESS\_ERRORS (*destination*, ATOMIC, MODIFY, WRITE\_CONTENTS)

OBJECT\_IS\_IN\_USE\_FOR\_MOVE (*destination*)

AUDIT\_FILE\_IS\_NOT\_ACTIVE (*source*)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_AUDIT)

PROCESS\_IS\_IN\_TRANSACTION

## 21.2.3 AUDIT\_FILE\_READ

```
AUDIT_FILE_READ (
    audit_file : Audit_file_designator
)
    records : Audit_file
```

AUDIT\_FILE\_READ reads the contents of the audit file *audit\_file*, returning the result as a sequence of auditing records in *records*.

#### Errors

ACCESS\_ERRORS (*audit\_file*, ATOMIC, READ, READ\_CONTENTS)

### 21.2.4 AUDIT\_GET\_CRITERIA

```
AUDIT_GET_CRITERIA (
    station      : Workstation_designator,
    criterion_type : Criterion_type
)
    criteria      : Criteria
```

AUDIT\_GET\_CRITERIA returns the set of criteria of the type given by *criterion\_type* that have been set for the workstation *station*. The returned set contains the event types that have been selected mapped to the return codes, mandatory labels, object designators, or user designators (depending on the *criterion\_type*) associated with each event type.

The set of criteria returned depends on the value of *criterion\_type*:

- GENERAL: the set of general criteria is returned.
- CONFIDENTIALITY\_DEPENDENT the set of confidentiality-dependent criteria is returned.
- INTEGRITY\_DEPENDENT: the set of integrity-dependent criteria is returned.
- OBJECT\_DEPENDENT: the set of object-dependent criteria is returned.
- USER\_DEPENDENT: the set of user-dependent criteria is returned.

#### Errors

OBJECT\_IS\_INACCESSIBLE (*station*, ATOMIC)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_AUDIT)

WORKSTATION\_IS\_UNKNOWN (*station*)

### 21.2.5 AUDIT\_RECORD\_WRITE

```
AUDIT_RECORD_WRITE (
    text      : String
)
```

AUDIT\_RECORD\_WRITE writes an information auditing record in the audit file *audit\_file* of the local workstation. The INFORMATION field of the auditing record is specified by *text*.

#### Errors

ACCESS\_ERRORS (*audit\_file*, ATOMIC, MODIFY, APPEND\_CONTENTS)

AUDIT\_FILE\_IS\_NOT\_ACTIVE (*audit\_file*)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_AUDIT\_INFORMATION\_LENGTH)

### 21.2.6 AUDIT\_REMOVE\_CRITERION

```
AUDIT_REMOVE_CRITERION (
    station      : Workstation_designator,
    criterion    : Removed_criterion
)
```

AUDIT\_REMOVE\_CRITERION removes the criterion *criterion* from the audit criteria of the workstation *station*.

For a selectable event type, all general selection criteria with that event type are removed regardless of the return code specified.

For a confidentiality-dependent criterion, events of the selected type performed on objects with the selected confidentiality label are no longer audited.

For an integrity-dependent criterion, events of the selected type performed on objects with the selected integrity label are no longer audited.

For an object-dependent criterion, events of the selected type performed on the selected object are no longer audited.

For a user-dependent criterion, events of the selected type performed on behalf of the selected user are no longer audited.

#### Errors

For confidentiality-dependent criterion:

```
CONFIDENTIALITY_CRITERION_IS_NOT_SELECTED (criterion)
CONFIDENTIALITY_LABEL_IS_INVALID (security label specified by criterion)
```

For event type:

```
EVENT_TYPE_IS_NOT_SELECTED (criterion)
```

For integrity-dependent criterion:

```
INTEGRITY_CRITERION_IS_NOT_SELECTED (criterion)
INTEGRITY_LABEL_IS_INVALID (security label specified by criterion)
```

For object-dependent criterion:

```
DISCRETIONARY_ACCESS_IS_NOT_GRANTED (object, ATOMIC)
OBJECT_CRITERION_IS_NOT_SELECTED (criterion)
```

For user-dependent criterion:

```
GROUP_IDENTIFIER_IS_INVALID (group identifier of criterion)
USER_IS_UNKNOWN (user specified by criterion)
USER_CRITERION_IS_NOT_SELECTED (criterion)
```

```
OBJECT_IS_INACCESSIBLE (station, ATOMIC)
```

```
PRIVILEGE_IS_NOT_GRANTED (PCTE_AUDIT)
```

```
WORKSTATION_IS_UNKNOWN (station)
```

### 21.2.7 AUDIT\_SELECTION\_CLEAR

```
AUDIT_SELECTION_CLEAR (
    station : Workstation_designator
)
```



AUDIT\_SELECTION\_CLEAR removes all selected audit criteria from the workstation *station*.

#### Errors

OBJECT\_IS\_INACCESSIBLE (*station*, ATOMIC)  
 PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_AUDIT)  
 WORKSTATION\_IS\_UNKNOWN (*station*)

### 21.2.8 AUDIT\_SWITCH\_OFF\_SELECTION

```
AUDIT_SWITCH_OFF_SELECTION (
    station : Workstation_designator
)
```

AUDIT\_SWITCH\_OFF\_SELECTION disables auditing on the workstation *station*. Events on *station* will no longer be audited, except for the event types that are always audited.

The current auditing selection criteria are maintained.

#### Errors

OBJECT\_IS\_INACCESSIBLE (*station*, ATOMIC)  
 PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_AUDIT)  
 WORKSTATION\_IS\_UNKNOWN (*station*)

### 21.2.9 AUDIT\_SWITCH\_ON\_SELECTION

```
AUDIT_SWITCH_ON_SELECTION (
    station : Workstation_designator
)
```

AUDIT\_SWITCH\_ON\_SELECTION enables auditing on the workstation *station*. Events on *station* will then be audited according to the current selection criteria. If auditing is already enabled, then the operation has no effect.

#### Errors

OBJECT\_IS\_INACCESSIBLE (*station*, ATOMIC)  
 PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_AUDIT)  
 WORKSTATION\_IS\_UNKNOWN (*station*)

### 21.2.10 AUDITING\_GET\_STATUS

```
AUDITING_GET_STATUS (
    station : Workstation_designator
)
    status : Audit_status
```

AUDITING\_GET\_STATUS returns ENABLED if auditing is currently enabled on the workstation *station*, and DISABLED otherwise.

#### Errors

OBJECT\_IS\_INACCESSIBLE (*station*, ATOMIC)

PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_AUDIT)  
 WORKSTATION\_IS\_UNKNOWN (*station*)

## 22 Accounting

### 22.1 Accounting concepts

#### 22.1.1 Consumers and accountable resources

Consumer\_identifier = Natural

Resource\_identifier = Natural

**sds** accounting:

**import object type** system-object, system-process, system-common\_root;

**import attribute type** system-number;

accounting\_directory: **child type of object with link**

known\_consumer\_group: (**navigate**) **existence link** (consumer\_identifier: **natural**) to consumer\_group;

known\_resource\_group: (**navigate**) **existence link** (resource\_identifier: **natural**) to resource\_group;

accounts\_of: **implicit link to common\_root reverse** accounts;

**end** accounting\_directory;

consumer\_group: **child type of object with link**

consumer\_process: (**navigate**) **non\_duplicated designation link** (number) to process;

**end** consumer\_group;

resource\_group: **child type of object with link**

resource\_group\_of: (**navigate**) **reference link** (number) to object **reverse** in\_resource\_group;

**end** resource\_group;

**extend object type** process **with link**

consumer\_identity: (**navigate**) **designation link to** consumer\_group;

**end** process;

**extend object type** object **with link**

in\_resource\_group: (**navigate**) **reference link to** resource\_group **reverse** resource\_group\_of;

**end** object;

**extend object type** common\_root **with link**

accounts: (**navigate**) **existence link to** accounting\_directory **reverse** accounts\_of;

**end** common\_root;

**end** accounting;

A consumer group is a group of consumer processes which are accounted together for their usage of accountable resources.

A resource group is a group of *accountable resources*, the usage of which are accounted together. *Accountable resources* are files, pipes, devices, static contexts, workstations, SDSs, and message queues. There is a "resource\_group\_of" link from the resource group to each of its accountable resources.

The accounting directory is an administrative object (see 9.1.2).

Each consumer group and each resource group has an associated consumer identifier or resource identifier respectively which identifies it uniquely within the PCTE installation and is used in the construction of accounting records. These identifiers are key attributes of the links from the accounting directory to the consumer group and resource group respectively.

A process may be associated with a consumer group, which is the destination of the "consumer\_identity" link. If a process is not associated with a consumer group, accounting is still effective for that process.

## 22.1.2 Accounting logs and accounting records

```
Accounting_log ::
  RECORDS : seq of Accounting_record
  represented by accounting_log

Accounting_record = Workstation_accounting_record
  | Static_context_accounting_record
  | Sds_accounting_record
  | Device_accounting_record
  | File_accounting_record
  | Pipe_accounting_record
  | Message_queue_accounting_record
  | Information_accounting_record

Basic_accounting_record ::
  SECURITY_USER      : Group_identifier
  ADOPTED_USER_GROUP : Group_identifier
  CONSUMER_GROUP     : [ Consumer_identifier ]
  RESOURCE_GROUP     : [ Resource_identifier ]
  START_TIME         : Time
  KIND               : Resource_kind

Resource_kind = WORKSTATION | STATIC_CONTEXT | SDS | DEVICE | FILE | PIPE |
MESSAGE_QUEUE | INFORMATION

Workstation_accounting_record :: Basic_accounting_record &&
-- KIND = WORKSTATION
  DURATION      : Float
  CPU_TIME      : Float
  SYS_TIME      : Float

Static_context_accounting_record :: Basic_accounting_record &&
-- KIND = STATIC_CONTEXT
  DURATION      : Float
  CPU_TIME      : Float
  SYS_TIME      : Float

Sds_accounting_record = Basic_accounting_record
-- KIND = SDS
```

```

Device_accounting_record :: Basic_accounting_record &&      -- KIND = DEVICE
    DURATION          : Float
    READ_COUNT        : Natural
    WRITE_COUNT       : Natural
    READ_SIZE         : Natural
    WRITE_SIZE        : Natural

File_accounting_record :: Basic_accounting_record &&        -- KIND = FILE
    DURATION          : Float
    READ_COUNT        : Natural
    WRITE_COUNT       : Natural
    READ_SIZE         : Natural
    WRITE_SIZE        : Natural

Pipe_accounting_record :: Basic_accounting_record &&        -- KIND = PIPE
    DURATION          : Float
    READ_COUNT        : Natural
    WRITE_COUNT       : Natural
    READ_SIZE         : Natural
    WRITE_SIZE        : Natural

Message_queue_accounting_record :: Basic_accounting_record &&
    -- KIND = MESSAGE_QUEUE
    OPERATION         : SEND | RECEIVE
    MESSAGE_SIZE       : Natural

Information_accounting_record :: Basic_accounting_record && -- KIND = INFORMATION
    INFORMATION        : String

sds accounting:
import object type system-workstation;

extend object type workstation with
link
    has_log: (navigate) reference link to accounting_log reverse is_log_for;
end workstation;

accounting_log: child type of object with
contents accounting_log;
link
    is_log_for: (navigate) reference link (number) to workstation reverse has_log;
end accounting_log;
end accounting;

```

An accounting log is an object associated with a workstation which is a server (see below). It has an "is\_log\_for" link to each associated workstation.

An accounting record is a record of accountable resource usage by a process. Each usage has a *start event* when the usage is deemed to start and an *end event* when it is deemed to be complete. The accounting record is written to the accounting log associated with the workstation which is a server for the accountable resource at the end event. The accountable resource usages are as follows.

- Use of the contents of a file, pipe, or device (KIND is FILE, PIPE or DEVICE respectively). The start event is when the process opens the contents (CONTENTS\_OPEN); the end event is when the process next closes the contents (CONTENTS\_CLOSE) or when the process terminates (PROCESS\_TERMINATE).

- Use of a static context or workstation associated with the process (KIND is STATIC\_CONTEXT or WORKSTATION respectively). The start event is when the process is started (PROCESS\_START or PROCESS\_CREATE\_AND\_START); the end event is when the process terminates (PROCESS\_TERMINATE).
- Use of an SDS in the working schema of the process (KIND is SDS). The start event is when the process is started (PROCESS\_START or PROCESS\_CREATE\_AND\_START) or when a working schema containing the SDS is set (PROCESS\_SET\_WORKING\_SCHEMA); the end event is when a new working schema is set (PROCESS\_SET\_WORKING\_SCHEMA) or the process terminates (PROCESS\_TERMINATE).
- Sending a message to a message queue or receiving a message from a message queue (KIND is MESSAGE\_QUEUE). The start and end events are the same: the sending or receiving of the message (MESSAGE\_SEND\_WAIT, MESSAGE\_SEND\_NO\_WAIT, MESSAGE\_RECEIVE\_WAIT, MESSAGE\_RECEIVE\_NO\_WAIT).
- Certain operations act as an end event followed by a start event for all started accounting resource usages by the calling process; they are PROCESS\_SET\_CONSUMER\_IDENTITY, PROCESS\_UNSET\_CONSUMER\_IDENTITY, PROCESS\_SET\_USER, and PROCESS\_ADOPT\_GROUP.
- Certain operations act as an end event for certain started accounting resource usages by the calling process; they are WORKSTATION\_DISCONNECT for accountable resources on volumes of the workstation; VOLUME\_UNMOUNT for accountable resources on the volume; PROCESS\_TERMINATE and ACTIVITY\_ABORT for started accounting resource usages by the process. ACCOUNTING\_OFF acts as an end event for all accountable resources on volumes of the workstation, for all processes.
- Certain events may be end events for started accounting resource usages by the calling process; they are failure of the execution site of the process, and the volume on which the accountable resource resides becoming inaccessible. In the case of static context and SDS resources on inaccessible volumes, whether such events are end events or not is implementation-defined.

When a resource is made accountable after its usage has started, or is removed from being accountable before its usage has ended, if such usage is recorded, and if so how, are implementation-defined.

If an accountable resource becomes inaccessible to the process, this counts as an end event for the usage of that resource.

A process may also write accounting records via ACCOUNTING\_RECORD\_WRITE (KIND is INFORMATION).

A workstation is a *server* for an accountable resource if the accountable resource resides on a volume mounted on a device controlled by the workstation, and the workstation is associated with an accounting log and accounting is enabled on the workstation.

The information in an accounting record depends on the kind of accountable resource involved. Each accounting record has a fixed part and a resource specific part. The fields of the accounting record are set as follows.

- Basic accounting record (fixed part):

- . SECURITY\_USER: the group identifier of the user identity of the process;
- . ADOPTED\_USER\_GROUP: the group identifier of the adopted user group of the process;
- . CONSUMER\_GROUP: the exact identifier of the consumer group of the process;
- . RESOURCE\_GROUP: the exact identifier of the resource group of the accountable resource;
- . START\_TIME: the time by the system clock at the start event of the usage of the accountable resource;
- . DURATION: the duration of the usage of the accountable resource, in seconds, from the start event to the end event;
- . INFORMATION: free for use by tools writing accounting records into the accounting log via ACCOUNTING\_RECORD\_WRITE; absent in other accounting records.
- Workstation accounting record and static context accounting record:
  - . CPU\_TIME: the consumption of processor time in seconds by the process during the usage of the workstation or static context;
  - . SYS\_TIME: the consumption of system time in seconds by the process during the usage of the workstation or static context.
- Device accounting record, File accounting record, or Pipe accounting record:
  - . READ\_COUNT: number of read operations by the process from the device, file, or pipe during the usage;
  - . WRITE\_COUNT: number of write operations by the process to the device, file, or pipe during the usage;
  - . READ\_SIZE: total size in octets of data read by the process from the device, file, or pipe during the usage;
  - . WRITE\_SIZE: total size in octets of data written by the process to the device, file, or pipe during the usage.

A read operation for device accounting purposes is a CONTENTS\_READ only. A write operation for device accounting purposes is CONTENTS\_WRITE or CONTENTS\_TRUNCATE.

- Message queue accounting record:

- . OPERATION: whether the usage was to send or to receive a message;
- . MESSAGE\_SIZE: the size in octets of the message sent or received.

The structure of the accounting log is implementation-defined.

If, when writing to an accounting log, the write fails because the accounting log is unavailable, then the operation which caused the accountable event to occur waits until an accounting log is made available. The means by which the accounting log unavailability is notified to the operators of the PCTE installation is implementation-defined. When an end event occurs, the completed accounting record is not lost. If there is an abnormal closedown of the workstation

before the end event of the usage of accountable resources, the extent to which the completion of such usages are recorded in further accounting records is implementation-defined.

#### NOTES

- 1 It is intended that accounting logs are persistent across workstation failures, and that modifications to an accounting log are not subject to the transaction rollback mechanism (i.e. updates to an accounting log can never be discarded).
- 2 The start of accountable usage for a resource may be when it is defined as an accountable resource, and the end of accountable usage may be when it is defined to be no longer an accountable resource. While the optimum implementation is to use the start and end of accountability of the resource as the start and end events of usage for resources which are in use at the time of change, it is recognized that this may cause unnecessary complication and inefficiency. An implementation must therefore specify how this situation is handled.
- 3 No constraints on the label of the accounting log are enforced by the system when it writes accounting records, although a bitwise write occurs. Whether this bitwise write, when it gives rise to a covert channel, is audited or not, is implementation-defined.

## 22.2 Accounting administration operations

### 22.2.1 ACCOUNTING\_LOG\_COPY\_AND\_RESET

```
ACCOUNTING_LOG_COPY_AND_RESET (
    source_log      : Accounting_log_designator
    destination_log : Accounting_log_designator
)
```

ACCOUNTING\_LOG\_COPY\_AND\_RESET appends the contents of the accounting log *source\_log* to the contents of the accounting log *destination\_log*. The contents of *source\_log* is then reset, i.e. the contents of *source\_log* is now empty.

There is no loss of accounting records.

A write lock of the default mode is obtained on *source\_log*.

#### Errors

ACCESS\_ERRORS(*source\_log*, ATOMIC, MODIFY, (READ\_CONTENTS, WRITE\_CONTENTS))

ACCESS\_ERRORS(*destination\_log*, ATOMIC, MODIFY, APPEND\_CONTENTS)

ACCOUNTING\_LOG\_IS\_NOT\_ACTIVE(*source\_log*)

OBJECT\_IS\_IN\_USE\_FOR\_MOVE(*destination\_log*)

PROCESS\_IS\_IN\_TRANSACTION

### 22.2.2 ACCOUNTING\_LOG\_READ

```
ACCOUNTING_LOG_READ (
    log      : Accounting_log_designator
)
    records : Accounting_log
```

ACCOUNTING\_LOG\_READ returns the sequence of accounting records in the accounting log *log*.



**Errors**

ACCESS\_ERRORS (*log*, ATOMIC, READ, READ\_CONTENTS)

**22.2.3 ACCOUNTING\_OFF**

```
ACCOUNTING_OFF (
    station : Workstation_designator
)
```

ACCOUNTING\_OFF disables the accounting mechanism on the workstation *station*.

The "has\_log" link from the object *station* (if there is one) and its reverse "is\_log\_for" link are deleted.

Write locks of the default mode are obtained on the deleted link.

**Errors**

ACCESS\_ERRORS (accounting log of *station*, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (*station*, ATOMIC, MODIFY, WRITE\_LINKS)

WORKSTATION\_IS\_UNKNOWN (*station*)

**22.2.4 ACCOUNTING\_ON**

```
ACCOUNTING_ON (
    log      : Accounting_log_designator,
    station  : Workstation_designator
)
```

ACCOUNTING\_ON enables the accounting mechanism on the workstation *station* with accounting log *log*.

A "has\_log" link, reversed by an "is\_log\_for" link, is created from *station* to *log*.

Write locks of the default mode are obtained on the created links.

**Errors**

ACCESS\_ERRORS (*log*, ATOMIC, MODIFY, APPEND\_CONTENTS)

ACCESS\_ERRORS (*station*, ATOMIC, MODIFY, APPEND\_LINKS)

LINK\_EXISTS (*station*, "has\_log" link)

OBJECT\_IS\_NOT\_ON\_ADMINISTRATION\_VOLUME (*log*, *station*)

WORKSTATION\_IS\_UNKNOWN (*station*)

NOTE - The error LINK\_EXISTS indicates that accounting was already enabled.

**22.2.5 ACCOUNTING\_RECORD\_WRITE**

```
ACCOUNTING_RECORD_WRITE (
    log          : Accounting_log_designator,
    information : String
)
```

ACCOUNTING\_RECORD\_WRITE appends a basic accounting record to the accounting log *log*.

The fields of the accounting record are set as follows:

- SECURITY\_USER is set to the group identifier of the destination of the "user\_identity" link of the calling process;
- ADOPTED\_USER\_GROUP is set to the group identifier of the adopted user group of the calling process;
- CONSUMER\_GROUP is set to the exact identifier of the destination of the "consumer\_identity" link of the calling process;
- RESOURCE\_GROUP is set to null;
- KIND is set to INFORMATION;
- START\_TIME is set to the current system time;
- INFORMATION is set to the parameter *information*.

#### Errors

ACCESS\_ERRORS (*log*, ATOMIC, MODIFY, APPEND\_CONTENTS)

ACCOUNTING\_LOG\_IS\_NOT\_ACTIVE (*log*)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_ACCOUNT\_INFORMATION\_LENGTH)

### 22.2.6 CONSUMER\_GROUP\_INITIALIZE

```
CONSUMER_GROUP_INITIALIZE (
    group      : Consumer_group_designator
)
    identifier : Consumer_identifier
```

CONSUMER\_GROUP\_INITIALIZE establishes the object *group* as a known consumer group. A "consumer\_group" link is created from the accounting directory to *group*. The key of this link is set to a unique value; which is guaranteed never to be re-used as a consumer group identifier, and is returned in *identifier*.

Write locks of the default mode are obtained on the created links.

#### Errors

ACCESS\_ERRORS (*group*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

ACCESS\_ERRORS (the accounting directory, ATOMIC, MODIFY, APPEND\_LINKS)

CONSUMER\_GROUP\_IS\_KNOWN (*group*)

### 22.2.7 CONSUMER\_GROUP\_REMOVE

```
CONSUMER_GROUP_REMOVE (
    group : Consumer_group_designator
)
```

CONSUMER\_GROUP\_REMOVE removes the consumer group *group* from the set of known consumer groups.

No process must currently use the "consumer\_identity" associated with that consumer group (i.e. the object *group* must not have any "consumer\_process" links).

The "known\_consumer\_group" link from the accounting directory to *group* is deleted.

If it is the only existence link to the object *group* and if there are no composition links to *group*, then *group* is also deleted. In that case, the "object\_on\_volume" link to *group* from the volume on which *group* was residing is also deleted.

Write locks of the default mode are obtained on the deleted links and on *group* if it is deleted.

#### Errors

ACCESS\_ERRORS (*group*, ATOMIC, CHANGE, WRITE\_IMPLICIT)

ACCESS\_ERRORS (the accounting directory, ATOMIC, MODIFY, WRITE\_LINKS)

If conditions hold for the deletion of the *group*:

ACCESS\_ERRORS (*group*, COMPOSITE, MODIFY, DELETE)

CONSUMER\_GROUP\_IS\_IN\_USE (*group*)

CONSUMER\_GROUP\_IS\_UNKNOWN (*group*)

OBJECT\_HAS\_LINKS\_PREVENTING\_DELETION (*group*)

OBJECT\_IS\_IN\_USE\_FOR\_DELETE (*group*)

### 22.2.8 RESOURCE\_GROUP\_ADD\_OBJECT

```
RESOURCE_GROUP_ADD_OBJECT (
    object : Object_designator,
    group  : Resource_group_designator
)
```

RESOURCE\_GROUP\_ADD\_OBJECT defines the object *object* to be an accountable resource.

An "in\_resource\_group" link is created from *object* to *group*. Its reverse "resource\_group\_of" link is keyed by the next unused value (see 23.1.2.7).

Write locks of the default mode are obtained on the created links.

#### Errors

ACCESS\_ERRORS (*object*, ATOMIC, MODIFY, APPEND\_LINKS)

ACCESS\_ERRORS (*group*, ATOMIC, MODIFY, APPEND\_LINKS)

OBJECT\_IS\_ALREADY\_IN\_RESOURCE\_GROUP (*object*, *group*)

OBJECT\_IS\_NOT\_ACCOUNTABLE\_RESOURCE (*object*)

RESOURCE\_GROUP\_IS\_UNKNOWN (*group*)

### 22.2.9 RESOURCE\_GROUP\_INITIALIZE

```
RESOURCE_GROUP_INITIALIZE (
    group      : Resource_group_designator
)
    identifier : Resource_identifier
```

RESOURCE\_GROUP\_INITIALIZE establishes the object *group* as a known resource group. A "resource\_group" link is created from the accounting directory to *group*. The key of this link is

set to a unique value, which is guaranteed never to be re-used as a resource group identifier, and is returned in *identifier*.

Write locks of the default mode are obtained on the created links.

#### Errors

ACCESS\_ERRORS (*group*, ATOMIC, CHANGE, APPEND\_IMPLICIT)

ACCESS\_ERRORS (the accounting directory, ATOMIC, MODIFY, APPEND\_LINKS)

RESOURCE\_GROUP\_IS\_KNOWN (*group*)

### 22.2.10 RESOURCE\_GROUP\_REMOVE

```
RESOURCE_GROUP_REMOVE (
    group : Resource_group_designator
)
```

RESOURCE\_GROUP\_REMOVE removes the resource group *group* from the set of known resource groups.

*group* must have no "resource\_group\_of" links to accountable resources.

The "known\_resource\_group" link from the accounting directory to *group* is deleted. If it is the only existence link to *group* and if there are no composition links to *group* then *group* is also deleted. In that case, the "object\_on\_volume" link to *group* from the volume on which the *group* was residing is also deleted.

Write locks of the default mode are obtained on the deleted links and on *group* if it is deleted.

#### Errors

ACCESS\_ERRORS (*group*, ATOMIC, CHANGE, WRITE\_IMPLICIT)

ACCESS\_ERRORS (the accounting directory, ATOMIC, MODIFY, WRITE\_LINKS)

If conditions hold for the deletion of the *group*:

ACCESS\_ERRORS (*group*, COMPOSITE, MODIFY, DELETE)

OBJECT\_HAS\_LINKS\_PREVENTING\_DELETION (*group*)

OBJECT\_IS\_IN\_USE\_FOR\_DELETE (*group*)

RESOURCE\_GROUP\_IS\_UNKNOWN (*group*)

### 22.2.11 RESOURCE\_GROUP\_REMOVE\_OBJECT

```
RESOURCE_GROUP_REMOVE_OBJECT (
    object : Object_designator,
    group : Resource_group_designator
)
```

RESOURCE\_GROUP\_REMOVE\_OBJECT removes the object *object* as an accountable resource from the resource group *group*.

The "resource\_group\_of" link and the "in\_resource\_group" reverse link between *object* and *group* are deleted.

Write locks of the default mode are obtained on the deleted links.

**Errors**

ACCESS\_ERRORS (*object*, ATOMIC, MODIFY, WRITE\_LINKS)

ACCESS\_ERRORS (*group*, ATOMIC, MODIFY, WRITE\_LINKS)

OBJECT\_IS\_NOT\_IN\_RESOURCE\_GROUP (*object*, *group*)

RESOURCE\_GROUP\_IS\_UNKNOWN (*group*)

**22.3 Consumer identity operations****22.3.1 PROCESS\_SET\_CONSUMER\_IDENTITY**

```
PROCESS_SET_CONSUMER_IDENTITY (
    group : Consumer_group_designator
)
```

PROCESS\_SET\_CONSUMER\_IDENTITY sets the consumer identity of the calling process, by creating a "consumer\_identity" link from the calling process to *group* and a complementary "consumer\_process" link from *group* to the calling process.

If the calling process already has a "consumer\_identity" link, that link and its complementary "consumer\_process" link are deleted.

**Errors**

ACCESS\_ERRORS (*group*, ATOMIC, SYSTEM\_ACCESS)

CONSUMER\_GROUP\_IS\_UNKNOWN (*group*)

DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED (*group*, ATOMIC, EXPLOIT\_CONSUMER\_IDENTITY)

If *process* is the calling process:

VOLUME\_IS\_FULL (calling process)

**22.3.2 PROCESS\_UNSET\_CONSUMER\_IDENTITY**

```
PROCESS_UNSET_CONSUMER_IDENTITY (
)
```

PROCESS\_UNSET\_CONSUMER\_IDENTITY suppresses the consumer identity of the calling process by deleting the "consumer\_identity" link, if any, from the calling process and a complementary "consumer\_process" link. If the calling process has no "consumer\_identity" link, the operation has no effect.

**Errors**

None.

## 23 Common binding features

### 23.1 Mapping of types

#### 23.1.1 Mapping of predefined PCTE datatypes

*Predefined PCTE datatypes* are the datatypes used as or to form the types of parameters and results of operations defined in this part of ISO/IEC 13719 which are not constructed from other PCTE datatypes. They are: Boolean, Natural, Integer, Float, Time, Text, and Octet. In order to define the possible values of these PCTE datatypes in a way which allows sensible binding decisions to be made, use is made of ISO/IEC 11404 which defines a comprehensive set of *LI datatypes* in abstract terms as sets of *values* and of associated *characterizing operations*. A language binding must define a mapping from these LI datatypes to binding language datatypes, which must ensure that all the characterizing operations are supported. In most cases the binding language supports them, but if not then the binding must supply them separately.

##### 23.1.1.1 Boolean values

The PCTE datatype Boolean is mapped to the primitive LI datatype boolean. This has 2 values, true and false; it is unordered. The characterizing operations are Equal, Not, And, Or.

##### 23.1.1.2 Integer values

The PCTE datatype Integer is mapped to the primitive LI datatype integer or to a generated LI datatype integer range (lowerbound .. upperbound) with appropriate bounds.

*integer* is a primitive LI datatype comprising the mathematical integers (positive, negative, and zero); its characterizing operations are Equal, Add, Multiply, Negate, NonNegative, and InOrder.

*range* is a subtype generator which creates a subtype of an ordered LI datatype within given bounds. The characterizing operations of the subtype are the same as those of the parent type.

The bounds are required to satisfy:

- upperbound = MAX\_INTEGER\_ATTRIBUTE;
- lowerbound = MIN\_INTEGER\_ATTRIBUTE.

##### 23.1.1.3 Natural values

The PCTE datatype Natural is mapped to a generated LI datatype integer range (0 .. upperbound) with appropriate upper bound.

The characterizing operations are as for the LI datatype integer except for Negate (and NonNegative which is not required as it is always true).

The upper bound is required to satisfy:

- upperbound = MAX\_NATURAL\_ATTRIBUTE.

#### 23.1.1.4 Float values

The PCTE datatype Float is mapped to a primitive LI datatype real (radix, factor), where radix and factor are integers with  $\text{radix} > 1$ , or to a generated LI datatype real (radix, factor) range (lowerbound .. upperbound). float (radix, factor) is a subset of the mathematical datatype of real numbers with precision of at least  $\text{radix} - \text{factor}$ . The characterizing operations are Equal, Add, Multiply, Negate, Reciprocal, and InOrder (and Promote which is not required as there is no PCTE datatype corresponding to the LI datatype rational).

The values radix, factor, lowerbound, and upperbound must satisfy the following:

- upperbound = MAX\_FLOAT\_ATTRIBUTE;
- lowerbound = MIN\_FLOAT\_ATTRIBUTE;
- $\text{radix} - \text{factor} \leq 10 - \text{MAX\_DIGITS\_FLOAT\_ATTRIBUTE}$ ;
- the smallest positive and negative numbers representable are  $\pm \text{SMALLEST\_FLOAT\_ATTRIBUTE}$ .

#### 23.1.1.5 Time values

The PCTE datatype Time is mapped to a primitive LI datatype time (second) or time (second, radix, factor), where radix and factor are integers with  $\text{radix} > 1$ , or to a subtype time (second) range (lowerbound .. upperbound) or time (second, radix, factor) range (lowerbound .. upperbound) with appropriate bounds. This is a datatype representing moments in time to a resolution of 1 second, or to a fraction of a second defined by  $\text{radix} - \text{factor}$ . The characterizing operations are Equal, InOrder, Difference, Extend (to a more precise resolution), and Round (to a less precise resolution).

The binding mapping must respect the limits:

- factor  $\geq 0$  (resolution at worst 1 second);
- upperbound = MAX\_TIME\_ATTRIBUTE;
- lowerbound = MIN\_TIME\_ATTRIBUTE.

#### 23.1.1.6 Octet, character, and text values

The PCTE datatype Octet is mapped to the LI datatype octet = new integer range (0 .. 255). The characterizing operation is Equal (and Select and Replace (from array) which are not required as there is no PCTE datatype corresponding to bit).

Octet values have no intrinsic graphical representation. When a graphical representation is required, the graphical representation of the PCTE datatype Character is used.

The PCTE datatype Character comprises the human-readable characters of one or more character sets selected by the PCTE implementation. In a character set, a single character may be represented by a single byte or by more than one byte.

The PCTE Datatype Text comprises sequences of characters. Characters of more than one character set may exist in a single text value. The method of identifying the character set of a given character is implementation-defined.



## NOTES

1 By the definition of the PCTE datatype Character, an octet may be part of a character of a multi-byte character set. Therefore, it may occur that an octet which is identical to an octet associated with a character of a single-byte character set is part of a character of a multi-byte character set. Even if such an octet is identical to an octet associated with a special character (e.g. '/', '\$', '#', '.' in a pathname), a PCTE implementation should not interpret the octet as such a special character.

2 For the definitions of the terms 'octet' and 'character set' see ISO/IEC 10646-1. For the definition of the term 'byte' see ISO/IEC 2022.

**23.1.1.7 Token values**

The PCTE datatype Token is used only as the field type of a single anonymous field of a record type, called a *private PCTE datatype*, as e.g. in:

User\_defined\_message\_type :: Token

This ensures that the values of type User\_defined\_message\_type are distinct from the values of all other types.

Except for designators and nominators (for which see 23.1.2), each private PCTE datatype is mapped to a distinct LI datatype new private (*length*), where private is an LI datatype defined by

type private (*length*: NaturalNumber) = new array (1.. *length*) of (bit)

For each such type *length* is a binding-defined positive integer.

**23.1.1.8 Enumeration values**

*Enumerated PCTE datatypes* are unions of VDM-SL enumeration types, each comprising a single enumerated value. An enumerated PCTE datatype:

VALUE1 | VALUE2 | ...

is mapped to an LI enumerated datatype enumerated (value1, value2, ...) with corresponding values. The characterizing operation is Equal (and InOrder and Successor which are not required as an enumerated PCTE datatype is unordered).

**23.1.2 Mapping of designators and nominators**

This clause defines the constraints on the PCTE datatypes used in bindings for referring to objects, attributes, links, and types, and types in SDS.

These datatypes are object designators, attribute designators, link designators, type nominators, type nominators in SDS, and actual keys, when used as or in parameters or results of operations defined in clauses 9 to 22; the corresponding binding types are called object references, attribute references, link references, type references, type names in SDS, and keys respectively. Object references, attribute references, link references, and type references are binding-defined datatypes supported by operations defined in 23.2 to 23.4. Type names in SDS and keys are text values with an internal syntax, defined in the BSI metasyntactic notation.

For all designators and nominators, except link designators used in the creation of links, the entity designated must always exist. The process of identifying the entity from a reference is called *evaluation* of the reference; this is implicitly performed by all operations in clauses 9 to

22 for unevaluated references. This process can give rise to error conditions, which are defined in 23.1.2.1 to 23.1.2.4.

NOTE - The mapping of the PCTE datatypes used in the bindings is summarized in table 11, with the operations used to create values of the types. Text creation is binding-defined.

**Table 11 - Mapping of PCTE datatypes to common binding datatypes**

| PCTE datatype         | Binding datatype    | Created by                                                      |
|-----------------------|---------------------|-----------------------------------------------------------------|
| object designator     | object reference    | OBJECT_REFERENCE_SET_ABSOLUTE,<br>OBJECT_REFERENCE_SET_RELATIVE |
| attribute designator  | attribute reference | TYPE_REFERENCE_SET                                              |
| link designator       | link reference      | LINK_REFERENCE_SET                                              |
| type nominator        | type reference      | TYPE_REFERENCE_SET                                              |
| type nominator in SDS | type name in SDS    | text creation                                                   |
| actual key            | key                 | text creation                                                   |

### 23.1.2.1 References

$X\_reference = Internal\_X\_reference \mid External\_X\_reference$

External\_X\_reference ::

NAME : X\_name

EVALUABILITY : Boolean

Evaluation\_point = NOW | FIRST\_USE | EVERY\_USE

Evaluation\_status = INTERNAL | EXTERNAL

Reference\_equality = EQUAL\_REFS | UNEQUAL\_REFS | EXTERNAL\_REFS

References are an abstract datatype characterized by the operations of 23.2 to 23.4; the above VDM-SL type definition of an X\_reference, where 'X' is 'object', 'attribute', 'link', or 'type', is for expository purposes only and need not be mapped explicitly in a binding.

References provide two ways of designating an object, attribute, link, or type: by a name (an *external reference*), and by a handle (an *internal reference*). The *evaluation status* of a reference is external or internal accordingly. An object, attribute, link, or type is accessed from a reference by *evaluating* it. The syntax of external references and the structure of internal references are defined in 23.1.2.2 to 23.1.2.5 inclusive. An external reference is evaluated by first *pre-evaluating* it to give an internal reference; pre-evaluation in general evaluates the reference as far as possible in the absence of any other information. The internal reference is then evaluated in the context of any other required information to complete the process.

The evaluability applies only to external references; it is **true** if the reference is to be converted to an internal reference when next pre-evaluated by an operation of clauses 9 to 22 and **false** otherwise. Evaluation points are used as parameters of operations returning references to indicate the evaluation status and evaluability required: NOW indicates an internal reference,

FIRST\_USE an external reference with evaluability **true**; and EVERY\_USE an external reference with evaluability **false**.

The evaluation of a reference takes place during the successful execution of an operation of clauses 9 to 22 of which the reference is a parameter. Operations in clause 23 do not as a rule evaluate their reference parameters, though in some cases (where stated) external reference parameters are pre-evaluated.

References returned by any of the operations in clauses 9 to 22 are always internal.

### 23.1.2.2 Object references

Internal\_object\_reference :: Token

An object reference identifies an object. The syntax of object names (also called *pathnames*) is as follows:

pathname = referenced object name, [ '/', relative pathname ] | ['\$current\_object', '/' ],  
relative pathname;

relative pathname = link name, {'/', link name};

referenced object name = '\$', key string value | alias;

For link references see 23.1.2.4.

Pre-evaluation of an external object reference is the same as evaluation, and evaluation of an internal object reference is a null process.

A relative pathname specifies a chain of links starting from a given origin object; the first link is specified by the origin object and the first link name; the second by the destination of the first link and the second link name, and so on. Finally the relative pathname specifies the destination object of the final link.

A pathname with no relative pathname specifies the same object as the referenced object name. A pathname with a relative pathname specifies the final destination object given by the object specified by the referenced object name and the relative pathname, as just described.

A pathname which consists only of a relative pathname is equivalent to a pathname starting from the current object and following the specified relative pathname, as just described.

A referenced object name of the first form specifies the destination of the "referenced\_object" link from the calling process with the key given by the key string value. The key is a string which is the referenced object name. If the link to reference object is omitted from an external object designator, the default reference object is ".".

For practical purposes, *aliases* are provided for the most commonly used referenced objects; see table 12.

**Table 12 - Aliases of referenced objects**

| Alias | Key of referenced object | Meaning                                                                                                                                                                                  |
|-------|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| "\$"  | "self"                   | The current process object                                                                                                                                                               |
| "#"   | "static_context"         | The static context of the current process                                                                                                                                                |
| "_"   | "common_root"            | The common root of the PCTE installation                                                                                                                                                 |
| "~"   | "home_object"            | An object conventionally associated with each user, called "home". The type of home is not predefined.                                                                                   |
| "."   | "current_object"         | An object conventionally chosen for the interpretation of a pathname without a starting referenced object name (see above) and providing the conventional notion of a current directory. |

When an object reference is evaluated, the constituent pathname, if any, is evaluated. The evaluation of the pathname involves the evaluation of the link names in the pathname.

The visible types are those that are visible at the time of calling an operation. Thus even if an internal object reference is used in an operation, the visible types are those that are visible when the operation is called rather than when the object reference is evaluated.

Evaluation of an external object reference *reference* may give rise to the following errors, which can therefore occur in any operation which has an object designator as parameter or result.

ACCESS\_ERRORS (object identified by *reference*, ATOMIC, READ, NAVIGATE)

For each link name *link* in the relative pathname:

ACCESS\_ERRORS (origin object of *link*, ATOMIC, SYSTEM\_ACCESS)

LINK\_DESTINATION\_DOES\_NOT\_EXIST (*link*)

LINK\_DESTINATION\_IS\_NOT\_VISIBLE (*link*)

USAGE\_MODE\_ON\_LINK\_TYPE\_WOULD\_BE\_VIOLATED (origin object of *link*, *link*, NAVIGATE)

Errors arising from evaluation of *link* (see 23.1.2.4)

LINK\_DESTINATION\_DOES\_NOT\_EXIST ("referenced\_object" link from the calling process identified by the referenced object name of *reference*)

REFERENCE\_CANNOT\_BE\_ALLOCATED

REFERENCE\_NAME\_IS\_INVALID (*reference*)

REFERENCED\_OBJECT\_IS\_UNSET (*reference*)

Any use of an object reference *reference* as parameter of an operation may additionally raise the following errors, whatever its evaluation status.

OBJECT\_IS\_OF\_WRONG\_TYPE (*reference*)

OBJECT\_REFERENCE\_IS\_INVALID (*reference*)

OBJECT\_REFERENCE\_IS\_UNSET (*reference*)

### 23.1.2.3 Attribute references

Internal\_attribute\_reference :: Internal\_type\_reference

An attribute reference identifies an attribute relative to a given object or link. It may therefore identify different attributes depending on the object or link. The syntax of attribute names is as follows:

attribute name = type name;

Pre-evaluation of an external attribute reference consists in evaluating the type name (see 23.1.2.5) as an attribute type name to give an internal type reference. Evaluation of an internal attribute reference consists in identifying the attribute of the given object or link with that attribute type; as no two attributes of an object or link may have the same attribute type, evaluation is unambiguous.

Evaluation of an attribute reference *reference* may give rise to the following errors:

ATTRIBUTE\_TYPE\_IS\_NOT\_VISIBLE (*reference*)

If the operation does not delete the attribute

ATTRIBUTE\_TYPE\_OF\_LINK\_TYPE\_IS\_NOT\_APPLIED (*reference*)

ATTRIBUTE\_TYPE\_OF\_OBJECT\_TYPE\_IS\_NOT\_APPLIED (*reference*)

Errors arising from evaluation of the attribute type reference (see 23.1.2.5).

### 23.1.2.4 Link references

Internal\_link\_reference ::

ACTUAL\_LINK\_TYPE : [Internal\_type\_reference]

KEY : [Key]

An internal link type reference must contain an actual link type or a key (or both). A key is a text value obeying the syntax of a key (see 23.1.2.7).

A link reference identifies a link in the context of its origin. It may therefore identify different links depending on the origin. The syntax of link names is as follows:

link name = cardinality one link name | cardinality many link name;

cardinality one link name = ':', type name;

cardinality many link name = key, ':', [ type name ] | key;

The *canonical form* of a link name is with a type identifier for link type name, and no '+', '++', or '-' key attribute values in the key.

The second form of cardinality many link name is allowed only if the key consists of a single key attribute value, or if the rightmost key attribute value of the key does not obey the syntax of a type name.

The pre-evaluation of an external link reference is as follows. If the link name contains a type name, that type name is evaluated as a link type name (see 23.1.2.5) to give the actual link type. If not, the internal link reference has no actual link type and the determination of the actual link type is deferred. For an external reference with a cardinality one link name, the internal link reference has no key; for an external link reference with a cardinality many link name, the key of the internal link reference is the same as the key of the link name.

The evaluation of an internal link reference is as follows:

- If the link reference does not contain an actual link type, the preferred link type of the origin is used as actual link type.
- If there is no key, the link reference identifies the cardinality one link with the given origin and the actual link type.
- If there is a key, it is evaluated in the context of the actual link type to yield an actual key, as described in 23.1.2.7. The identified link is the cardinality many link with the given origin, the actual link type, and the actual key.

Pre-evaluation of a link reference *reference* may give rise to the following errors.

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_KEY\_VALUE)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_KEY\_SIZE)

LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_LINK\_NAME\_SIZE)

Errors arising from evaluation of the link type reference (see 23.1.2.5)

Evaluation of a link reference *reference* in the context of the origin *origin* may additionally give rise to the following errors, whatever the evaluation status of *reference*.

If any '+' or '++' key attribute values are evaluated:

DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED (*origin*, ATOMIC, READ\_LINKS)

If any '-' key attribute values are evaluated:

DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED (*origin*, ATOMIC, READ\_ATTRIBUTES)

KEY\_IS\_BAD (*origin*, *reference*)

LINK\_DOES\_NOT\_EXIST (*origin*, *reference*)

LINK\_TYPE\_IS\_NOT\_APPLIED\_TO\_OBJECT\_TYPE (object type of *origin*, link type reference of *reference*)

If the link type name is not supplied:

PREFERENCE\_DOES\_NOT\_EXIST (*origin*, *reference*)

Errors arising from evaluation of the preferred link type reference (see 23.1.2.5).

### 23.1.2.5 Type references

Internal\_type\_reference :: Token

A type reference identifies a type. The syntax of a type name is as follows:

type name = local name | full type name | type identifier;

local name = name;

full type name = sds name, '-', local name;

sds name = name;

type identifier = '\_', string;

Pre-evaluation of an external type reference is the same as evaluation, and evaluation of an internal type reference is a null process.

A full type name identifies a type with the given local name in the SDS specified by the SDS name which is a member of the sequence of SDS names in the current working schema. A type name which is just a local name identifies the type in the current working schema. If derived from a local name, the evaluation of the type reference yields the first type in working schema (in the sequence of SDS names in the current working schema) with that local name as its local name.

A type identifier is a string with first character '\_'; the syntax and interpretation of the rest of the string are implementation-defined. The value of the "type\_identifier" attribute of a "type" object (see 10.1.2) is the corresponding type identifier without the initial '\_' character.

If a type name or a link name is returned by an operation rather than a type reference or a link reference, the type name or link name is returned as a type name if the type is visible and named in the current working schema, and as a type identifier otherwise. When a type name is returned it is the local name of the first named associated type in SDS in the sequence of SDS names in the working schema, provided that this local name does not occur earlier in the sequence of SDS names for another type. In the latter case, the full type name, i.e. prefixed by an SDS name, of the first associated type in SDS in the sequence of SDS names in the working schema is returned.

The use of a type identifier as or as part of an input parameter is allowed if the type is visible, or if the predefined program group PCTE\_CONFIGURATION is effective for the calling process; creating objects and links, getting, setting, and resetting attributes (including the working schema operations of 10.4), and converting objects by means of types which are not visible are invalid even for the PCTE\_CONFIGURATION program group.

Evaluation of a type reference *reference* may give rise to the following errors:

If PCTE\_CONFIGURATION is not effective for the calling process:

ATTRIBUTE\_TYPE\_IS\_NOT\_VISIBLE (*reference*)

ENUMERAL\_TYPE\_IS\_NOT\_VISIBLE (*reference*)

LINK\_TYPE\_IS\_NOT\_VISIBLE (*reference*)

OBJECT\_TYPE\_IS\_NOT\_VISIBLE (*reference*)

If PCTE\_CONFIGURATION is effective for the calling process:

OPERATION\_IS\_NOT\_ALLOWED\_ON\_TYPE (*reference*)

If *reference* contains a type identifier *identifier*, the following implementation-defined error may be raised:

TYPE\_IDENTIFIER\_IS\_INVALID (*identifier*)

If *reference* is a full type name:

SDS\_IS\_NOT\_IN\_WORKING\_SCHEMA (SDS name of *reference*)

SDS\_IS\_UNKNOWN (SDS name of *reference*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (SDS name of *reference*)

TYPE\_IS\_NOT\_DESCENDANT (*reference*, expected type)

TYPE\_IS\_OF\_WRONG\_KIND (*reference*)

Any use of an internal type reference *reference* as a parameter of an operation may additionally raise the following error:

TYPE\_REFERENCE\_IS\_INVALID (*reference*)

NOTES



- 1 The ability of the program group PCTE\_CONFIGURATION to identify types which are not part of the working schema of the calling process, is intended to be used to remove garbage from the object base, e.g. "type" objects associated with types which are no longer in existence.
- 2 Internal type references are independent of the working schema, and remain valid over changes to the working schema.

### 23.1.2.6 Type names in SDS

A type nominator in SDS in clauses 9 to 22 corresponds to a type name in SDS in a language binding. A type name in SDS is a text value with the following syntax.

type name in sds = local name | type identifier;

A type name in SDS identifies a type in SDS in a given SDS. A type nominator in SDS returned by an operation is returned as a local name if the type is named in the relevant SDS, otherwise as a type identifier.

Evaluation of a type name in SDS *name* may give rise to the following errors:

TYPE\_IDENTIFIER\_USAGE\_IS\_INVALID (*name*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (given SDS, *name*)

### 23.1.2.7 Keys

A value of type Actual\_key in clauses 9 to 22 corresponds to a key in a language binding. A key is a text value with the following syntax.

key = key attribute value, {'.', key attribute value};

key attribute value = key natural value | key string value | key nil value;

key string value = key first character, {key character};

key first character = key character - ('\$' | '#' | '~' | '\_' | '+');

key character = character - (',' | '-' | '/');

key natural value = '0' | nonzero digit, {digit} | highest used value | next unused value;

nonzero digit = '1' | '2' | '3' | '4' | '5' | '6' | '7' | '8' | '9';

digit = '0' | nonzero digit;

highest used value = '+';

next unused value = '++';

key nil value = '-';

A key identifies a link in the context of a given origin object and an actual link type. It is evaluated to give an *actual key* according to the rules given below; the actual key is a sequence of key values (strings or naturals), and identifies the link with given type and origin object, and with that sequence of values of its key attributes. The length of a key is limited to MAX\_KEY\_SIZE and the values of key attribute values are limited to MAX\_KEY\_VALUE and MAX\_KEY\_SIZE (see 24.1).

The key attribute values in the key are evaluated in order giving key values of the actual key as follows.

- A key string value gives that string.

- A key natural value of the first or second form gives the natural number which it represents in the usual decimal representation.
- The highest used value '+' gives the greatest key attribute value, if any, in that position in the sequence of key values among all links which have the same origin and the same key prefix (the preceding sequence of key values evaluated according to these rules). If there are no such links, the value of '+' is zero.
- The next unused value '++' gives the value of '+' plus one when the actual key is to be used as the key of a link created by the operation, and the value of '+' in other cases.
- The key nil value '-' is undefined if the origin object of the link has no preferred link key, or if its preferred link type is not the given actual link type. Otherwise it gives the value of the key attribute in the same position in the preferred link key of the origin object. If the preferred link key attribute value is '+' or '++', it is evaluated as described above.
- If fewer key values are present than the number of key attribute types of the given link type, then the number is effectively made up by adding further '-' key attribute values, except that if the origin object of the link has no preferred link key, or if its preferred link type is not the given actual link type, a missing key attribute value corresponding to a string key attribute gives the empty string.

An actual key returned as or as part of the result of an operation has the form of a key with no '+', '++', or '-' key attribute values.

#### NOTES

- 1 Although a key must contain at least one key attribute value, the effect of an empty key can be obtained by using a key '-' consisting of a single key nil value.
- 2 All names are valid key string values.

### 23.1.3 Mapping of other values

#### 23.1.3.1 Security labels

A value of type `Security_label` in clauses 9 to 22 corresponds to a security label in a language binding. A security label is a text value with the following syntax.

```
security_label = class name | conjunction | disjunction | '*';
conjunction = unit, {space, 'AND', space, unit};
disjunction = unit, {space, 'OR', space, unit};
unit = class name | '(', security_label, ')';
class name = name;
space = (* space character *);
```

A class name corresponds to the mandatory class designator of the security or integrity class which is the destination of a "known\_mandatory\_class" link from the mandatory directory with that class name as "name" key attribute value. The units in a conjunction or disjunction correspond to the security labels of the UNITS field in some order; all orders are equivalent.

Class names are interpreted as confidentiality class names in confidentiality labels and as integrity class names in integrity labels. The security label value '\*' is valid only as the high

label of a range and represents `MAXIMUM_LABEL` (see 20.1.5). A null label is represented by an empty text value.

Evaluation of a security label can give rise to the following error:

`CLASS_NAME_IS_INVALID` (*name*)

### 23.1.3.2 Names

A value of type `Name` is represented by a string having the syntax of a name according to the following rules.

`name` = `name first character`, { `name character` };

`name first character` = `name character` - ('/', '\$', '#', '-', '\_', ':', '~', '\*', '(');

`name character` = `character` - ('/', '\$', '#', '-', '~', ')');

The following synonyms for national characters (i.e. outside the ISO 646 invariant subset) are permitted in names: /S for \$, /T for ~, and /H for #.

Names are used as SDS names and local names of types within SDSs, and as class names. The syntax is the same as for identifiers in DDL (see B.7).

### 23.1.3.3 Other compound types

There are no constraints on the mapping of values of other VDM-SL compound types: sequences, sets, optional types, product types, record types, union types, and map types. Bindings may specify rules for special values, e.g. for a set containing all possible elements.

## 23.2 Object reference operations

### 23.2.1 OBJECT\_REFERENCE\_COPY

```
OBJECT_REFERENCE_COPY (
  reference      : Object_reference,
  point         : Evaluation_point
)
  new_reference : Object_reference
```

`OBJECT_REFERENCE_COPY` returns a new object reference *new\_reference* designating the same object as *reference*. The evaluation status and evaluability of *new\_reference* are specified by *point* (see 23.1.2.1).

If *reference* is external and *point* is NOW, then *reference* is pre-evaluated.

#### Errors

`EVALUATION_STATUS_IS_INCONSISTENT_WITH_EVALUATION_POINT` (*reference*, *point*)

**23.2.2 OBJECT\_REFERENCE\_GET\_EVALUATION\_POINT**

```

OBJECT_REFERENCE_GET_EVALUATION_POINT (
    reference : Object_reference
)
    point : Evaluation_point

```

OBJECT\_REFERENCE\_GET\_EVALUATION\_POINT returns an evaluation point indicating the evaluation status and evaluability of the object reference *reference*, as defined in 23.1.2.1.

**Errors**

None.

**23.2.3 OBJECT\_REFERENCE\_GET\_PATH**

```

OBJECT_REFERENCE_GET_PATH (
    reference : Object_reference
)
    pathname : Pathname

```

OBJECT\_REFERENCE\_GET\_PATH returns the pathname of the external object reference *reference*.

**Errors**

OBJECT\_REFERENCE\_IS\_INTERNAL (*reference*)

**23.2.4 OBJECT\_REFERENCE\_GET\_STATUS**

```

OBJECT_REFERENCE_GET_STATUS (
    reference : Object_reference
)
    status : Evaluation_status

```

OBJECT\_REFERENCE\_GET\_STATUS returns the evaluation status of the object reference *reference*.

**Errors**

OBJECT\_REFERENCE\_IS\_UNSET (*reference*)

**23.2.5 OBJECT\_REFERENCE\_SET\_ABSOLUTE**

```

OBJECT_REFERENCE_SET_ABSOLUTE (
    pathname : Pathname,
    point : Evaluation_point
)
    new_reference : Object_reference

```

OBJECT\_REFERENCE\_SET\_ABSOLUTE creates a new object reference *new\_reference* from a pathname *pathname* and an evaluation point *point*. The evaluation status and evaluability of *new\_reference* are specified by *point* (see 23.1.2.1).

**Errors**

PATHNAME\_SYNTAX\_IS\_WRONG (*pathname*)

**23.2.6 OBJECT\_REFERENCE\_SET\_RELATIVE**

```
OBJECT_REFERENCE_SET_RELATIVE (
    reference      : Object_reference,
    pathname      : Relative_pathname,
    point         : Evaluation_point
)
    new_reference : Object_reference
```

OBJECT\_REFERENCE\_SET\_RELATIVE creates a new object reference from an existing object reference *reference*, a relative pathname *pathname*, and an evaluation point *point*. The evaluation status and evaluability of *new\_reference* are specified by *point* (see 23.1.2.1).

If *reference* is external and *point* is NOW, then *reference* is pre-evaluated.

**Errors**

EVALUATION\_STATUS\_IS\_INCONSISTENT\_WITH\_EVALUATION\_POINT (*reference*, *point*)

If *point* is NOW:

REFERENCE\_CANNOT\_BE\_ALLOCATED

OBJECT\_REFERENCE\_IS\_INVALID (*reference*)

PATHNAME\_SYNTAX\_IS\_WRONG (*pathname*)

**23.2.7 OBJECT\_REFERENCE\_UNSET**

```
OBJECT_REFERENCE_UNSET (
    reference : Object_reference
)
```

OBJECT\_REFERENCE\_UNSET deletes the object reference *reference*, releasing any associated resources.

**Errors**

OBJECT\_REFERENCE\_IS\_UNSET (*reference*)

**23.2.8 OBJECT\_REFERENCES\_ARE\_EQUAL**

```
OBJECT_REFERENCES_ARE_EQUAL (
    first_reference : Object_reference,
    second_reference : Object_reference
)
    equal           : Reference_equality
```

OBJECT\_REFERENCES\_ARE\_EQUAL returns EQUAL\_REFS if both object references *first\_reference* and *second\_reference* are internal and designate the same object; UNEQUAL\_REFS if both object references are internal and designate different objects; and EXTERNAL\_REFS otherwise (i.e. if one or both are external). REFERENCES\_ARE\_EQUAL does not evaluate either object reference.

**Errors**OBJECT\_REFERENCE\_IS\_UNSET (*first\_reference*)OBJECT\_REFERENCE\_IS\_UNSET (*second\_reference*)**23.3 Link reference operations****23.3.1 LINK\_REFERENCE\_COPY**

```

LINK_REFERENCE_COPY (
    reference      : Link_reference
    point         : Evaluation_point
)
    new_reference : Link_reference

```

LINK\_REFERENCE\_COPY returns a new link reference *new\_reference* which would identify the same link relative to a given object as *reference*. The evaluation status and evaluability of *new\_reference* are specified by *point* (see 23.1.2.1).

If *reference* is external and *point* is NOW, then *reference* is pre-evaluated.

**Errors**EVALUATION\_STATUS\_IS\_INCONSISTENT\_WITH\_EVALUATION\_POINT (*reference*, *point*)LINK\_REFERENCE\_IS\_UNSET (*reference*)**23.3.2 LINK\_REFERENCE\_GET\_EVALUATION\_POINT**

```

LINK_REFERENCE_GET_EVALUATION_POINT (
    reference : Link_reference
)
    point     : Evaluation_point

```

LINK\_REFERENCE\_GET\_EVALUATION\_POINT returns an evaluation point indicating the evaluation status and evaluability of the link reference *reference*, as defined in 23.1.2.1.

**Errors**LINK\_REFERENCE\_IS\_UNSET (*reference*)**23.3.3 LINK\_REFERENCE\_GET\_KEY**

```

LINK_REFERENCE_GET_KEY (
    reference : Link_reference
)
    key      : Key

```

LINK\_REFERENCE\_GET\_KEY returns the key of the link reference *reference*.

The key *key* is the key before any evaluation.

**Errors**LINK\_REFERENCE\_IS\_UNSET (*reference*)

### 23.3.4 LINK\_REFERENCE\_GET\_KEY\_VALUE

```
LINK_REFERENCE_GET_KEY_VALUE (
    reference : Link_reference,
    index     : Natural
)
key_value : Natural | Text
```

LINK\_REFERENCE\_GET\_KEY\_VALUE returns the key value *key\_value* indexed by *index* of the link reference *reference*, if this key value exists.

The first key value of the key of *reference* has index 1, the second 2, and so on.

#### Errors

KEY\_VALUE\_DOES\_NOT\_EXIST (*reference*, *index*)

LINK\_REFERENCE\_IS\_UNSET (*reference*)

### 23.3.5 LINK\_REFERENCE\_GET\_NAME

```
LINK_REFERENCE_GET_NAME (
    reference : Link_reference
)
link_name : Link_name
```

LINK\_REFERENCE\_GET\_NAME returns the link name of the link reference *reference*.

If *reference* is external, then *reference* is pre-evaluated.

If the link type of *reference* is visible then the link type name of *link\_name* is the local name of the first type in working schema in the sequence of SDSs in the working schema, or if there is no local name, the link type name of *link\_name* is the full type name of the first type in working schema in the sequence of SDSs in the working schema.

#### Errors

LINK\_NAME\_IS\_TOO\_LONG\_IN\_CURRENT\_WORKING\_SCHEMA (link name of *reference*)

LINK\_REFERENCE\_IS\_UNSET (*reference*)

### 23.3.6 LINK\_REFERENCE\_GET\_STATUS

```
LINK_REFERENCE_GET_STATUS (
    reference : Link_reference
)
status      : Evaluation_status
```

LINK\_REFERENCE\_GET\_STATUS returns the evaluation status of the link reference *reference*.

#### Errors

LINK\_REFERENCE\_IS\_UNSET (*reference*)



### 23.3.7 LINK\_REFERENCE\_GET\_TYPE

```
LINK_REFERENCE_GET_TYPE (
    reference          : Link_reference
)
    type_reference    : Link_type_reference
```

LINK\_REFERENCE\_GET\_TYPE returns the link type reference of the link reference *reference*.

#### Errors

LINK\_REFERENCE\_IS\_UNSET (*reference*)

### 23.3.8 LINK\_REFERENCE\_SET

```
LINK_REFERENCE_SET (
    link_name          : Link_name | Type_reference | (Key * Type_reference)
    point              : Evaluation_point
)
    new_reference : Link_reference
```

LINK\_REFERENCE\_SET creates a new link reference *new\_reference* from a link name *link\_name*. The evaluation status and evaluability of *new\_reference* are specified by *point* (see 23.1.2.1).

If *link\_name* is provided as a link name (i.e. as a text value), then it must conform to the syntax rules defined in 23.1.2.4.

If *link\_name* is provided as a type reference *reference*, the link reference *new\_reference* designates a cardinality one link with *reference* as a link type reference.

If *link\_name* is provided as a key *key* and a type reference *reference*, the link reference *new\_reference* designates a cardinality many link with *key* as a key and *reference* as a link type reference.

#### Errors

KEY\_SYNTAX\_IS\_WRONG (*key*)

If *link\_name* is a type reference or a key and a type reference:

EVALUATION\_STATUS\_IS\_INCONSISTENT\_WITH\_EVALUATION\_POINT  
(*reference*, *point*)

If *link\_name* is a link name:

LINK\_NAME\_SYNTAX\_IS\_WRONG (*link\_name*)

### 23.3.9 LINK\_REFERENCE\_UNSET

```
LINK_REFERENCE_UNSET (
    reference : Link_reference
)
```

LINK\_REFERENCE\_UNSET deletes the link reference *reference*, releasing any associated resources.

**Errors**LINK\_REFERENCE\_IS\_UNSET (*reference*)**23.3.10 LINK\_REFERENCES\_ARE\_EQUAL**

```

LINK_REFERENCES_ARE_EQUAL (
    first_reference    : Link_reference,
    second_reference  : Link_reference
)
    equal              : Reference_equality

```

LINK\_REFERENCES\_ARE\_EQUAL returns EQUAL\_REFS if both link references *first\_reference* and *second\_reference* are internal, they have textually equal keys, and their link types are equal as defined by TYPE\_REFERENCES\_ARE\_EQUAL (see 23.4.8); UNEQUAL\_REFS if both references are internal, but do not satisfy the equality rules for EQUAL\_REFS; and EXTERNAL\_REFS otherwise (i.e. if one or both are external). LINK\_REFERENCES\_ARE\_EQUAL does not evaluate either link reference.

**Errors**LINK\_REFERENCE\_IS\_UNSET (*first\_reference*)LINK\_REFERENCE\_IS\_UNSET (*second\_reference*)**23.4 Type reference operations****23.4.1 TYPE\_REFERENCE\_COPY**

```

TYPE_REFERENCE_COPY (
    reference          : Type_reference,
    point              : Evaluation_point
)
    new_reference     : Type_reference

```

TYPE\_REFERENCE\_COPY returns a new type reference *new\_reference* identifying the same type as *reference*. The evaluation status and evaluability of *new\_reference* are specified by *point* (see 23.1.2.1).

If *reference* is external and *point* is NOW, then *reference* is pre-evaluated.

**Errors**EVALUATION\_STATUS\_IS\_INCONSISTENT\_WITH\_EVALUATION\_POINT (*reference*, *point*)TYPE\_REFERENCE\_IS\_UNSET (*reference*)**23.4.2 TYPE\_REFERENCE\_GET\_EVALUATION\_POINT**

```

TYPE_REFERENCE_GET_EVALUATION_POINT (
    reference          : Type_reference
)
    point              : Evaluation_point

```

TYPE\_REFERENCE\_GET\_EVALUATION\_POINT returns an evaluation point indicating the evaluation status and evaluability of the type reference *reference*, as defined in 23.1.2.1.

**Errors**

TYPE\_REFERENCE\_IS\_UNSET (*reference*)

**23.4.3 TYPE\_REFERENCE\_GET\_IDENTIFIER**

```
TYPE_REFERENCE_GET_IDENTIFIER (
    reference : Type_reference
)
    identifier : Type_name
```

TYPE\_REFERENCE\_GET\_IDENTIFIER returns the type identifier *identifier* of the type reference *reference*.

If *reference* is external, then *reference* is pre-evaluated.

**Errors**

TYPE\_IS\_NOT\_VISIBLE (*reference*)

TYPE\_REFERENCE\_IS\_UNSET (*reference*)

**23.4.4 TYPE\_REFERENCE\_GET\_NAME**

```
TYPE_REFERENCE_GET_NAME (
    sds : [ Sds_designator ],
    reference : Type_reference
)
    name : Type_name
```

TYPE\_REFERENCE\_GET\_NAME returns the type name *name* of the type reference *reference*.

If *sds* is not provided, *name* is the local name, full type name, or type identifier, according to the rules for returned names in 23.1.2.5.

If *sds* is provided, *name* is the local name of the associated type in SDS in the SDS *sds*.

If *reference* is external, then *reference* is pre-evaluated.

**Errors**

If *sds* is supplied:

ACCESS\_ERRORS (*sds*, ATOMIC, READ, READ\_LINKS)

SDS\_IS\_UNKNOWN (*sds*)

TYPE\_HAS\_NO\_LOCAL\_NAME (*sds*, *reference*)

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *reference*)

TYPE\_REFERENCE\_IS\_UNSET (*reference*)

**23.4.5 TYPE\_REFERENCE\_GET\_STATUS**

```
TYPE_REFERENCE_GET_STATUS (
    reference : Type_reference
)
    status : Evaluation_status
```

TYPE\_REFERENCE\_GET\_STATUS returns the evaluation status of the type reference *reference*.

**Errors**

TYPE\_REFERENCE\_IS\_UNSET (*reference*)

**23.4.6 TYPE\_REFERENCE\_SET**

```

TYPE_REFERENCE_SET (
    name      : Type_name,
    point     : Evaluation_point
)
    new_reference : Type_reference

```

TYPE\_REFERENCE\_SET creates a new type reference *new\_reference* from a type name *name* and an evaluation point *point*. The evaluation status and evaluability of *new\_reference* are specified by *point* (see 23.1.2.1).

**Errors**

TYPE\_IDENTIFIER\_IS\_INVALID (*name*)

If *point* is NOW:

TYPE\_IS\_NOT\_VISIBLE (*name*)

TYPE\_NAME\_IS\_INVALID (*name*)

**23.4.7 TYPE\_REFERENCE\_UNSET**

```

TYPE_REFERENCE_UNSET (
    reference : Type_reference
)

```

TYPE\_REFERENCE\_UNSET deletes the type reference *reference* releasing any associated resources.

**Errors**

TYPE\_REFERENCE\_IS\_UNSET (*reference*)

**23.4.8 TYPE\_REFERENCES\_ARE\_EQUAL**

```

TYPE_REFERENCES_ARE_EQUAL (
    first_reference : Type_reference,
    second_reference : Type_reference
)
    equal          : Reference_equality

```

TYPE\_REFERENCES\_ARE\_EQUAL returns EQUAL\_REFS if both type references *first\_reference* and *second\_reference* are internal and designate the same type; UNEQUAL\_REFS if both references are internal and designate different types; and EXTERNAL\_REFS otherwise (i.e. if one or both are external). TYPE\_REFERENCES\_ARE\_EQUAL does not evaluate either type reference.

**Errors**

TYPE\_REFERENCE\_IS\_UNSET (*first\_reference*)

TYPE\_REFERENCE\_IS\_UNSET (*second\_reference*)

## 24 Implementation limits

Implementations may impose limits on the range, size, or number of certain quantities. Any attempt to exceed these limits gives rise to an error condition in the operation concerned. These error conditions are defined in the appropriate operation definitions.

This section defines bounds on these implementation-defined limits. Each bound defines the most constrained value of the limit that an implementation may impose, and is therefore the limit that should be assumed for portability by a tool writer.

The limits fall into two categories:

- *installation-wide limits*, which must be the same for all workstations in a PCTE installation;
- *workstation-dependent limits*, which may vary for different workstations in a PCTE installation.

### 24.1 Bounds on installation-wide limits

**MAX\_ACCESS\_CONTROL\_LIST\_LENGTH:** The maximum number of entries which may appear in each of the atomic ACLs of an object or the default access control list of a process, as a natural; must be at least 64.

**MAX\_ACCOUNT\_DURATION, DELTA\_ACCOUNT\_DURATION:** Upper and lower limits respectively for duration values in an accounting record, as floats. **MAX\_ACCOUNT\_DURATION** must be at least 86400 seconds; **DELTA\_ACCOUNT\_DURATION** must be at most 1 second.

**MAX\_ACCOUNT\_INFORMATION\_LENGTH:** The maximum number of octets in the **INFORMATION** field of an accounting record, as a natural; must be at least 128.

**MAX\_AUDIT\_INFORMATION\_LENGTH:** The maximum number of octets in the **TEXT** part of an audit record, as a natural; must be at least 128.

**MAX\_DIGIT\_FLOAT\_ATTRIBUTE:** The precision of a float attribute value, in decimal digits, as a natural; must be at least 6.

**MAX\_FLOAT\_ATTRIBUTE, MIN\_FLOAT\_ATTRIBUTE:** Upper and lower limits respectively for float attribute values, as floats. **MAX\_FLOAT\_ATTRIBUTE** must be at least  $10^{32}$ . **MIN\_FLOAT\_ATTRIBUTE** must be negative and its absolute value must be at least that of **MAX\_FLOAT\_ATTRIBUTE**.

**MAX\_INTEGER\_ATTRIBUTE, MIN\_INTEGER\_ATTRIBUTE:** Upper and lower limits respectively for integer attribute values, as integers. **MAX\_INTEGER\_ATTRIBUTE** must be at least 2147483647; **MIN\_INTEGER\_ATTRIBUTE** must be negative, and its absolute value must be one greater than **MAX\_INTEGER\_ATTRIBUTE**.

**MAX\_KEY\_SIZE:** The maximum number of octets in a string key attribute value, as a natural; must be at least 127.

**MAX\_KEY\_VALUE:** The maximum natural value in a natural key attribute value; must be at least 32000.

**MAX\_LINK\_NAME\_SIZE:** The maximum number of octets in a link reference, as a natural; must be at least 191.

MAX\_MESSAGE\_SIZE: The maximum number of octets in the data a message, as a natural; must be at least 1024.

MAX\_NAME\_SIZE: The maximum number of octets in a name or enumerals type image, as a natural; must be at least 31.

MAX\_NATURAL\_ATTRIBUTE: The upper limit for a non-key natural attribute value, as a natural. Must be at least 2147483647.

MAX\_PRIORITY\_VALUE: The maximum priority value for a process, as a natural; must be at least 31.

MAX\_SECURITY\_GROUPS: The maximum number of security groups which may be initialized in a PCTE installation, as a natural; must be at least 32000.

MAX\_STRING\_ATTRIBUTE\_SIZE: The maximum number of octets in a string non-key attribute value, as a natural; must be at least 32000.

MAX\_TIME\_ATTRIBUTE: The latest time attribute value, as a time value; must be no earlier than 2044-12-31T24:00:00Z (24:00 on 31 December 2044 UTC).

MIN\_TIME\_ATTRIBUTE: The earliest time attribute value, as a time value; must be no later than 1980-01-01T00:00:00Z (00:00 on 1 January 1980 UTC).

SMALLEST\_FLOAT\_ATTRIBUTE: The lower limit on the absolute value of float attribute values, as a float; must be at greatest  $10^{-32}$ .

## 24.2 Bounds on workstation-dependent limits

MAX\_ACTIVITIES: The maximum number of activities on the workstation, as a natural; must be at least 256.

MAX\_ACTIVITIES\_PER\_PROCESS: The maximum number of activities for a process executing on the workstation, as a natural; must be at least 8. MAX\_ACTIVITIES\_PER\_PROCESS is also the maximum depth of nesting of activities on the workstation.

MAX\_FILE\_SIZE: The maximum size in octets of a file on the workstation, as a natural; must be at least 100,000,000.

MAX\_MESSAGE\_QUEUE\_SPACE: The maximum total space of a message queue on the workstation, as a natural; must be at least 32000.

MAX\_MOUNTED\_VOLUMES: The maximum number of volumes mounted on devices controlled by this workstation, as a natural; must be at least 16.

MAX\_OPEN\_OBJECTS: The maximum number of concurrently open objects on the workstation, as a natural; must be at least 512.

MAX\_OPEN\_OBJECTS\_PER\_PROCESS: The maximum number of concurrently open objects for a process executing on the workstation, as a natural; must be at least 16.

MAX\_PIPE\_SIZE: The maximum size in octets in a pipe on the workstation, as a natural; must be at least 4096.

**MAX\_PROCESSES:** The maximum number of processes running, stopped, or suspended on the workstation, as a natural; must be at least 64.

**MAX\_PROCESSES\_PER\_USER:** The maximum number of processes existing simultaneously and created by a user on the workstation, as a natural; must be at least 16.

**MAX\_SDS\_IN\_WORKING\_SCHEMA:** The maximum number of SDSs in a working schema on the workstation, as a natural; must be at least 32.

## 24.3 Limit operations

### 24.3.1 Datatypes for limit operations

Limit\_category = STANDARD | IMPLEMENTATION | REMAINING

Limit\_name = MAX\_ACCESS\_CONTROL\_LIST\_LENGTH | MAX\_ACCOUNT\_DURATION |  
 DELTA\_ACCOUNT\_DURATION | MAX\_ACCOUNT\_INFORMATION\_LENGTH |  
 MAX\_ACTIVITIES | MAX\_ACTIVITIES\_PER\_PROCESS |  
 MAX\_AUDIT\_INFORMATION\_LENGTH | MAX\_DIGIT\_FLOAT\_ATTRIBUTE |  
 MAX\_FILE\_SIZE | MAX\_FLOAT\_ATTRIBUTE | MIN\_FLOAT\_ATTRIBUTE |  
 MAX\_INTEGER\_ATTRIBUTE | MIN\_INTEGER\_ATTRIBUTE | MAX\_KEY\_SIZE |  
 MAX\_KEY\_VALUE | MAX\_LINK\_REFERENCE\_SIZE | MAX\_MESSAGE\_QUEUE\_SPACE |  
 MAX\_MESSAGE\_SIZE | MAX\_MOUNTED\_VOLUMES | MAX\_NAME\_SIZE |  
 MAX\_NATURAL\_ATTRIBUTE | MAX\_OPEN\_OBJECTS |  
 MAX\_OPEN\_OBJECTS\_PER\_PROCESS | MAX\_PIPE\_SIZE | MAX\_PRIORITY\_VALUE |  
 MAX\_PROCESSES | MAX\_PROCESSES\_PER\_USER |  
 MAX\_SDS\_IN\_WORKING\_SCHEMA | MAX\_SECURITY\_GROUPS |  
 MAX\_STRING\_ATTRIBUTE\_SIZE | MAX\_TIME\_ATTRIBUTE | MIN\_TIME\_ATTRIBUTE |  
 SMALLEST\_FLOAT\_ATTRIBUTE

Limit\_value = Natural | Integer | Float | Time

These datatypes are used in the operation LIMIT\_GET\_VALUE.

### 24.3.2 LIMIT\_GET\_VALUE

```
LIMIT_GET_VALUE (
  limit_category : Limit_category,
  limit_name     : Limit_name
)
  limit_value    : [ Limit_value ]
```

LIMIT\_GET\_VALUE returns for the limit named by *limit\_name* the limit bound given in 24.1 or 24.2, the limit imposed by the implementation, or (where relevant) the current available quantity of the resource in the PCTE installation or local workstation, according as *limit\_category* is STANDARD, IMPLEMENTATION, or REMAINING respectively. If the implementation does not impose a particular limit, then IMPLEMENTATION returns no value for that limit.

REMAINING returns a value for the following values of *limit\_name*, otherwise no value is returned: MAX\_SECURITY\_GROUPS, MAX\_ACTIVITIES, MAX\_MOUNTED\_VOLUMES, MAX\_OPEN\_OBJECTS, MAX\_PROCESSES.

#### Errors

None.



## Annex A (normative)

### VDM Specification Language for the Abstract Specification

#### A.1 Introduction

The Abstract Specification uses a very limited subset of VDM-SL, the Specification Language of the Vienna Development Method as defined in ISO/IEC DIS 13817. This annex defines the subset, explains the semantics of the subset informally, and defines concrete syntax for use in the Abstract Specification.

The definition of VDM-SL in ISO/IEC DIS 13817 is principally in terms of its abstract syntax, i.e. the bare structure of the language with all concrete syntactic details removed. It is permissible to use *any* concrete syntax which can be mapped to the abstract syntax, but ISO/IEC DIS 13817 does define two particular concrete syntaxes, called the mathematical and the ISO 646 syntaxes. As their names suggest, the former uses many mathematical and quasimathematical symbols, and is intended for typeset text; the latter uses only the ISO 646 character set and is intended for use with unsophisticated data preparation equipment and for information interchange. The two are virtually isomorphic down to the lexical level.

The VDM-SL subset used in this part of ISO/IEC 13719 uses the ISO 646 syntax, which for the subset is perfectly readable and avoids the need for special symbols. Use is also made of a convention, defined below, for different styles for different kinds of identifiers.

#### A.2 The VDM-SL subset

The subset consists of:

- type definitions, to define the types of the various conceptual entities;
- state definitions, to define the PCTE state;
- a simple form of operation definitions, to define operation headings;
- a simple form of function definition, to define auxiliary function headings;
- operation and function calls;
- identifiers;
- literals, to express values.

There are three small extensions. The first is of a purely syntactic nature; it has been found very useful for defining composite types as extensions of previously defined composite types (the '&&' notation). The second is used to relate the VDM-SL and the DDL definitions (the '**represented by**' notation). The third is to allow an operation to return more than result; it is simple in each case to map such an operation to an equivalent operation returning a single result of a composite type with the results of the original operation as its fields.

## A.2.1 Type definitions

### Syntax

type definition = identifier, '=', type expression | identifier, '::', [identifier, '&&'], field list,  
 ['**represented**', '**by**', identifier];  
 type expression = bracketed type expression | type name | quote type expression | set type  
 expression | map type expression | sequence type expression | union type expression | optional  
 type expression | product type expression;  
 bracketed type expression = '(', type expression, ')';  
 type name = identifier;  
 quote type expression = quote literal;  
 set type expression = '**set**', '**of**', type expression;  
 map type expression = '**map**', type expression, '**to**', type expression;  
 sequence type expression = '**seq**', '**of**', type expression | '**seq1**', '**of**', type expression;  
 union type expression = type expression, '|', type expression;  
 optional type expression = '[', type expression, ']';  
 product type expression = type expression, '\*', type expression;  
 field list = {field};  
 field = [identifier, ':'], type expression;

### Semantics

A *type definition* declares a *type*, i.e. a set of values, and associates it with an identifier. The declared type is defined as follows, according to the different kinds of type definition.

- $I = T$ . The type denoted by the identifier  $I$  is defined by the type expression  $T$ .
- $I :: I_1 : T_1 \ I_2 : T_2 \ \dots$ . The type denoted by  $I$  is a *composite type*: each value of the type is in effect an ordered set of values called *fields*, one from each of the field types  $T_1, T_2, \dots$  in that order. The field identifiers  $I_1, I_2, \dots$  are used in full VDM-SL to access the fields; in our subset they are descriptive only. The different notation is to emphasize that values from two different composite types are always distinct, even though they have the same field types and values: composite types use name equivalence, whereas other types use structural equivalence.

The form  $I :: I' \ \&\& \ I_1 : T_1 \ I_2 : T_2 \ \dots$  is short for  $I :: I_1' : T_1' \ I_2' : T_2' \ \dots \ I_1 : T_1 \ I_2 : T_2 \ \dots$  where  $I_1' : T_1' \ I_2' : T_2' \ \dots$  are the fields of the composite type  $I'$ .

The '**represented by**  $I'$ ', if present, indicates that the type is represented by the predefined DDL object type definition with local name  $I'$ . All the local names of object types in the predefined SDSs are different, so there is no ambiguity.

The various forms of type expression define types as follows.

- bracketed type expression ( $T$ ). This denotes the same type as the type expression  $T$ ; the brackets are used to override or emphasize the precedence of the type operators.
- type name  $I$ . This denotes the type associated with the identifier  $I$  in a type definition. The predefined PCTE types (which take the place of VDM-SL basic types) are defined in 23.1.1.

- quote type expression. This denotes a *quote type* containing a single *quote value* denoted by the same quote literal as the quote type expression. Quote values have no predefined properties except equality and inequality. Types containing quote values are built up as union types, e.g.:

Colour = RED | GREEN | BLUE

Result = Natural | ERROR

- set type expression: **set of** T. Values of this type are the finite subsets of values of type T.
- sequence type expression: **seq of** T, **seq1 of** T. Values of these types are finite ordered sequences of values of type T; in the first case including, and in the second case excluding, the empty sequence.
- map type expressions: **map** T1 **to** T2. A value of either of these types is a finite *map* from type T1 to type T2, i.e. an association with each of a finite subset of T1 (the *domain* of the map) of a value of T2. The element of T2 associated with an element  $x$  of the domain of a map  $M$  is denoted by  $M(x)$ .
- union type expression: T1 | T2 | ... . This denotes the union of the constituent types T1, T2, ... ; a value of a union type is a value of one of the constituent types (which must be disjoint).
- optional type expression: [T]. This denotes the same type as T | **nil**, where **nil** is the only value of a special anonymous type (the *nil type*). **nil** is used conventionally to stand for absence of a value.
- product type expression: T1 \* T2 \* ... . Values of a product type are ordered tuples of values from the constituent types T1, T2, ... . The distinction from a composite type is that product types use structural equivalence, and the components are not named.

NOTE - Optional types are sometimes used to indicate the preferred default value.

## A.2.2 State definitions

### Syntax

state definition = **'state'**, identifier, **'of'**, field list, **'end'**;

### Semantics

The *state definition*

```

state I of
  I1 : T1
  I2 : T2
  ...
end
```

defines the state to consist of variables I1, .. of types T1, ... respectively. In the VDM-SL subset the identifier I is descriptive only (in full VDM-SL it is used to identify the state from within a different module).

### A.2.3 Operation headings

#### Syntax

operation heading = identifier, '(', [typed parameter list], ')', [typed results];

typed parameter list = identifier, ':', type expression, {'', identifier, ':', type expression};

typed results = identifier, ':', type expression, {'', identifier, ':', type expression};

#### Semantics

The *operation heading*

```
Op (
  I1: T1,
  I2: T2,
  ...,
)
  R1 : T1'
  R2 : T2'
  ...
```

declares Op as the name of an operation with formal parameters I1 of type T1, I2 of type T2, etc., and with results R1 of type T1', R2 of type T2', etc. If R1 : T1' etc. is omitted, the operation has no result. Note that the ( ) are needed even if the operation has no parameters.

### A.2.4 Function headings

#### Syntax

function heading = identifier, '(', [typed parameter list], ')', identifier, ':', type expression;

#### Semantics

The *function heading*

```
Fn (
  I1 : T1,
  I2 : T2,
  ...
)
  R : Tr
```

declares Fn as the name of a function with formal parameters I1 of type T1, I2 of type T2, ..., and with result R of type Tr. Note that the ( ) are needed even if the function has no parameters.

### A.2.5 Operation and function calls

#### Syntax

operation call = identifier, '(', [expression, {'', expression}], ')';

function application = identifier, '(', [expression, {'', expression}], ')';

#### Semantics

The operation call or function application

Id (E1, E2, ... )

denotes a call to the operation or function with name Id, with actual parameters E1, E2, ... .

## A.2.6 Identifiers

### Syntax

identifier = plain letter , {plain letter | digit | '\_'};

plain letter = capital letter | small letter;

capital letter = 'A' | 'B' | 'C' | 'D' | 'E' | 'F' | 'G' | 'H' | 'I' | 'J' | 'K' | 'L' | 'M' | 'N' | 'O' | 'P' | 'Q' | 'R' | 'S' | 'T' | 'U' | 'V' | 'W' | 'X' | 'Y' | 'Z';

small letter = 'a' | 'b' | 'c' | 'd' | 'e' | 'f' | 'g' | 'h' | 'i' | 'j' | 'k' | 'l' | 'm' | 'n' | 'o' | 'p' | 'q' | 'r' | 's' | 't' | 'u' | 'v' | 'w' | 'x' | 'y' | 'z';

digit = '0' | '1' | '2' | '3' | '4' | '5' | '6' | '7' | '8' | '9';

### Semantics

Identifiers have no intrinsic meaning; they are used as names of various entities.

## A.2.7 Literals

### Syntax

symbolic literal = boolean literal | numeric literal | character literal | text literal | quote literal;

boolean literal = 'true' | 'false';

numeric literal = numeral, ['.', digit, {digit}], [exponent];

exponent = 'E', ['+' | '-'], numeral;

numeral = digit, {digit};

character literal = "'", character, "'";

text literal = '"', {character - '"'}, '"';

quote literal = capital letter, {'\_' | capital letter};

### Semantics

Literals denote values of basic types.

Boolean literals denote the corresponding truth values.

Numeric literals denote rational numbers in the usual decimal notation, using '.' as the decimal point. An exponent  $E+n$  (or  $En$ ) or  $E-n$  denotes multiplication by  $10^n$  or  $10^{-n}$  respectively.

A character literal denotes a graphic character. A text literal denotes a sequence of characters, i.e. a value of type seq of char; as usual, the double quote character " is denoted by two successive double quote characters : "".

A quote literal denotes the only value of a quote type, and also the quote type itself.

## A.3 Conventions for identifiers and keywords

Identifiers are used in the VDM-SL subset to name entities of the kinds shown below. The conventions for these identifiers used in the Abstract Specification are shown. Note that two

identifiers cannot be distinguished, strictly speaking, just by the use of italics or not, as the ISO 646 syntax only recognizes one set of letters (capitals and small letters).

- Types: small letters with capital initial, e.g. Object, Type\_in\_sds.
- Fields of composite types: capital letters, e.g. DIRECT\_COMPONENTS.
- States: there is only one, PCTE\_Installation.
- State variables: capital letters, e.g. OBJECT\_BASE.
- Operations: capital letters, e.g. OBJECT\_GET\_ATTRIBUTE
- Operation parameters and results: small italic letters: *queue*, *next\_message*.

The same conventions are used in English text, except that names of types, fields, states, state variables, and values are converted to conventional English phrases (usually by replacing underscores by spaces and capital by small letters, but not always: PCTE installation, type in SDS).

In the ISO 646 syntax keywords are essentially reserved; there is a way of using a keyword as an identifier, but this is not used in the Abstract Specification. The keywords in the subset are:

|              |              |             |              |            |             |            |
|--------------|--------------|-------------|--------------|------------|-------------|------------|
| <b>end</b>   | <b>false</b> | <b>map</b>  | <b>of</b>    | <b>seq</b> | <b>seq1</b> | <b>set</b> |
| <b>state</b> | <b>to</b>    | <b>true</b> | <b>value</b> |            |             |            |

## Annex B (normative)

### The Data Definition Language (DDL)

This annex defines the PCTE Data Definition Language (DDL). DDL is used to define SDSs and the type definitions within them, and so serves for the definition of the schema of a PCTE installation. This definition serves as a language standard for DDL.

#### B.1 SDSs and clauses

##### Syntax

```
DDL definition = sds section, {sds section};

sds section =
    'sds', sds name, ':',
    {clause, ';'},
    'end', sds name, ';;';

clause =
    type importation | object type declaration | object type extension |
    attribute type declaration | link type declaration | link type extension |
    enumeration type declaration;

type importation =
    'import', import type, global name, [ 'as', local name ], [ type mode declaration ],
    { ':', global name, [ 'as', local name ], [ type mode declaration ] };

import type = 'object', 'type' | 'attribute', 'type' | 'link', 'type' | 'enumeration', 'type';
```

##### Meaning

A *DDL definition* defines a number of SDSs containing types in SDS, and the corresponding types.

All the *SDS sections* in the schema declaration with a particular SDS name define an SDS (schema definition set). This SDS has the common SDS name, and the set of types in SDS defined by the SDS sections as explained below.

The SDS section in which a clause occurs, and the SDS to which it contributes, are called the *current* SDS section and SDS, respectively.

The *type importation*

```
import import_type sds_name-local_name_1 as local_name_2 [type_mode_declaration]
```

defines a type in SDS in the current SDS. This type in SDS is a copy of the type in SDS denoted by local\_name\_1 in the SDS denoted by sds\_name, with the same type identifier, except that the creation or importation date is set to the value of the system clock at the time of importation. It is denoted by local name local\_name\_2 in the current SDS; if local\_name\_2 is not given, the default is local\_name\_1. Except for an enumeration type, the type mode declaration defines the usage and export modes of the new type in SDS; they may not contain any access values not in the export mode of the imported type. The default for both modes is the export mode of the imported type. The maximum mode is set to the export mode of the imported type.



A multiple type importation:

```
import import_type sds_name1-local_name1 [as new_local_name1] [type_mode_declaration1],
      sds_name2-local_name2 [as new_local_name2] [type_mode_declaration2],
      ...,
      sds_namen-local_namen [as new_local_namen] [ type_mode_declarationn ]
```

is equivalent to  $n$  single type importations:

```
import import_type sds_name1-local_name1 [as new_local_name1 ]
      [ type_mode_declaration1 ];
import import_type sds_name2-local_name2 [as new_local_name2 ]
      [ type_mode_declaration2 ];
      ...
import import_type sds_namen-local_namen [as new_local_namen ]
      [ type_mode_declarationn ]
```

Every type name used in an SDS must be declared in the same SDS, in a type importation, a type declaration, or type extension. Except for destination object types and reverse link types in link type declarations, all type names must be declared before use.

If a type mode declaration is omitted, all the definition mode values are set for the imported type.

## B.2 Object types

### Syntax

```
object type declaration =
local name, ':', [ type mode declaration ], [ child, 'type', 'of', object type list ], [ 'with',
[ 'contents', contents type indication, ':' ],
[ 'attribute',
      attribute indication list, ':' ],
[ 'link',
      link indication list, ':' ],
[ 'component',
      component indication list, ':' ],
'end', local name ];
```

```
object type extension =
'extend', 'object', 'type', local name, 'with',
[ 'attribute',
      attribute indication list, ':' ],
[ 'link',
      link indication list, ':' ],
[ 'component',
      component indication list, ':' ],
'end', local name;
```

contents type indication = **'file'** | **'pipe'** | **'device'** | **'audit\_file'** | **'accounting\_log'**;

attribute indication list = attribute indication list item, { ';', attribute indication list item };

attribute indication list item = attribute type name | attribute type declaration;

link indication list = link indication list item, { ';', link indication list item };

link indication list item = link type name | link type declaration;

component indication list = component indication list item, { ';', component indication list item };

component indication list item = link type name | link type declaration;

## Constraints

The '**child type of**' clause may be omitted only for the object type "object" (see 9.1.1).

The local name after '**end**' in an object type declaration or object type extension, if present, must be the same as the first local name of that object type declaration or object type extension.

In an object type declaration the local name must be distinct from the local names of all other object types, and of all attribute types and link types, defined in the same SDS as the object type declaration.

In an object type extension the local name must be the name of an object type introduced earlier in the SDS by an object type declaration or a type importation.

Each attribute type name in an attribute indication list must be the local name of an attribute type introduced earlier in the specification by an attribute type declaration or a type importation.

Each link type name in a link indication list must be the local name of a non-composition link type introduced earlier in the specification by an object or link type declaration or a type importation.

Each link type name in a component indication list must be the local name of a composition link type introduced earlier in the specification by an object or link type declaration or a type importation.

The type mode declaration must define either **protected** or **create** for each of the usage mode and the export mode.

All the attribute types and link types in the list must be different.

If any parents of an object type have contents, then they must all have the same contents type (there may be other parents with no contents). The child type inherits the common contents type, or if none of its parents has contents, neither has the child.

## Meaning

An *object type declaration* defines an object type, and an object type in SDS in the current SDS with the local name within that SDS. The new object type has the following characteristics (see 8.3.1).

- The contents type is as defined by the contents type indication. If a contents type indication is given, it defines the contents type of the object type as FILE, PIPE, DEVICE, AUDIT\_FILE, or ACCOUNTING\_LOG respectively (see 8.3.1). The contents type indication is used only for the predefined types "file", "pipe", "device", "audit\_file", and "accounting\_log"; user-defined types always inherit contents type (or lack of it ) from their parents.
- The parent types are the object types defined by the object type list after '**child type of**'; the object type is added to the child types of all its parent types. The object type has no child types initially.

The new object type in SDS has the following characteristics (see 8.4.1).

- The direct outgoing link types in SDS are all those defined by the link indication list and component indication list (see below).
- The direct attribute types in SDS are all those defined by the attribute indication list (see below)

- The direct component object types in SDS are all those defined as the destination object types of the composition links defined in the component indication list.
- The usage and export modes are set by the type mode declaration, if present (see B.6); the default is usage mode and export mode both set to CREATE.
- The maximum usage mode is set to CREATE.

An *object type extension* extends the object type in SDS with the same local name in the current SDS section, by adding further outgoing link types, attribute types, and component object types, defined as for an object type declaration.

An attribute indication list defines the following set of attribute types.

- For each attribute indication list item which is an attribute type name, the attribute type with that local name in the current SDS.
- For each attribute indication list item which is an attribute type declaration, all the attribute types defined by that attribute type declaration (see clause B.3).

A link indication list or a component indication list defines the following set of link types.

- For each link indication list item or component indication list item which is a link type name, the link type with that local name in the current SDS.
- For each link indication list item or component indication list item which is a link type declaration, the link type defined by that link type declaration (see clause B.4).

### B.3 Attribute types

#### Syntax

```

attribute type declaration =
    local name, {',' local name}, ':', [ type mode declaration ], [ 'non_duplicated' ],
    value type indication, [ ':=' , initial value ];

value type indication = 'integer' | 'natural' | 'boolean' | 'time' | 'float' | 'string' |
    'enumeration', enumeration type name | enumeration type indication;

enumeration type indication = 'enumeration', '(', basic enumeration,
    {',' basic enumeration}, ')';

basic enumeration = enumeration image | enumeration subrange;

enumeration image = identifier | ' " ', {character}, ' " ' ;

enumeration subrange = attribute type name, 'range', enumeration image, '..',
    enumeration image;

initial value =
    ['+' | '-'], digit, {digit}                (* integer *)
    | digit, {digit}                            (* natural *)
    | 'true' | 'false'                          (* boolean *)
    | year, '-', month, '-', day, ['T', hour, ':', minute, ':', second], 'Z'    (* time *)
    | ['+' | '-'], digit, {digit}, ['.', digit, {digit}], ['E', ['+' | '-'], digit, {digit}] (* float *)
    | '"', {character}, '"'                    (* string *)
    | enumeration image;                      (* enumeration value type *)

day = digit, digit;

month = digit, digit;

```

year = [digit, digit], digit, digit;

hour = digit, digit;

minute = digit, digit;

second = digit, digit;

### Constraints

All the local names of an attribute type declaration must be distinct from the local names of all other attribute types, and of all object types and link types, defined in the current SDS.

The initial value, if any, in an attribute type declaration must denote a value of the value type defined by the value type indication.

In a basic enumeration which is an enumeration subrange, the two enumeration images must identify different enumerals of the enumeration value type of the enumeration attribute type denoted by the attribute type name. In that enumeration value type the enumeral type identified by the second enumeration image must not precede that identified by the first.

The type mode declaration must define either **protected** or one or both of **read** and **write** for each of the usage mode and the export mode.

### Meaning

An *attribute type declaration* defines an attribute type, and an attribute type in SDS in the current SDS with the local name within that SDS. The new attribute type has the following characteristics (see 8.3.2).

- The value type is as given by the value type indication.
- The initial value of the attribute type is defined by the initial value of the attribute type declaration, if present, as follows.
  - . An initial value `['+' | '-'], digit, {digit}` denotes an integer in the conventional decimal representation.
  - . An initial value `'true'` or `'false'` denotes the logical value TRUE or FALSE respectively.
  - . An initial value year, '-', month, '-', day, `['T' hour, ':', minute, ':', second]` 'Z' denotes the time consisting of the given year, month, day, and the time of day given by hour, minute, second (if present) or otherwise 00:00:00, in Coordinated Universal Time (UCT). When omitted, the first two digits of the year are taken to be '19', thus '88-11-09Z' denotes the same time as '1988-11-09T00:00:00Z'.
  - . An initial value `['+' | '-'], digit, {digit}, ['.', digit, {digit}], ['E', ['+' | '-'], digit, {digit}]` denotes a rational number in the conventional decimal representation, where '.' represents the decimal point and E a decimal exponent. Leading and trailing zeros are permitted.
  - . An initial value `'"', {character}, '"'` denotes the string value consisting of the sequence of characters, without the delimiting `'"'` characters, except that the character `'"'` within the string is denoted by `'\"'`.
  - . An initial value which is an enumeration image denotes the corresponding enumeral type.
- The duplication is NON\_DUPLICATED if the keyword **'non\_duplicated'** appears, and otherwise DUPLICATED.

The new attribute type in SDS has the following properties (see 8.4.2).

- The usage and export modes are set by the type mode declaration, if present (see clause B.6); the default is usage mode and export mode both set to READ and WRITE.
- The maximum usage mode is set to READ and WRITE.

An enumeration type indication defines an enumeration value type consisting of the concatenation of the sequences of enumerals defined by the constituent basic enumerations, as follows. An enumeration image which is delimited by ' " " ' is interpreted as for a string attribute initial value.

- If the basic enumeration is an enumeration image, then a sequence containing the single numeral type with that image in the current SDS.
- If the basic enumeration is an enumeration subrange, then the sequence containing the enumerals between and including those identified by the two enumeration images, in the ordering of the enumeration value type of the attribute type.

## B.4 Link types

### Syntax

link type declaration =  
 local name, ':', [ type mode declaration ], [ '**exclusive**' ], [ '**non\_duplicated**' ],  
 [ stability name ], category name, '**link**', [ cardinality range ], [ key list ],  
 [ '**to**', object type list ], [ '**reverse**', link type name ], [ '**with**',  
 '**attribute**',  
 attribute indication list, ';',  
 '**end**', local name ];

link type extension =  
 '**extend**', '**link**', '**type**', local name, [ '**to**', object type list ], [ '**with**',  
 '**attribute**',  
 attribute indication list, ';',  
 '**end**', local name ];

category name = [ '**composition**' ] | '**existence**' | '**reference**' | '**implicit**' | '**designation**';

cardinality range = '[', [ lower bound ], '..', [ upper bound ], ']';

lower bound = digit, { digit };

upper bound = digit, { digit };

stability name = '**atomic**', '**stable**' | '**composite**', '**stable**';

key list = '(', attribute indication list, ')';

### Constraints

In a cardinality range:

- The lower bound must be greater than or equal to zero.
- The upper bound must be greater than zero.
- The lower bound must be less than or equal to the upper bound.
- Links of category name **implicit** or **existence** must have a lower bound equal to zero.

If the upper bound of a link type is greater than one, and the link type has category name **implicit**, then the key list must consist of a single attribute type name. If the upper bound of a link type is greater than one, then the key list must not be omitted.

The link type denoted by the link type name after '**reverse**', if present, must be a direct outgoing link type of each destination object type of the link type, and must have the link type as its reverse. See 8.3.3 for constraints on the properties of the reverse link type.

The type mode declaration must define either **protected** or a combination of one or more of **create**, **navigate**, and **delete** for each of the usage mode and the export mode.

### Meaning

A *link type declaration* defines a link type, and a link type in SDS in the current SDS with the local name within that SDS. The new link type has the following properties (see 8.3.3).

- Category: COMPOSITION, EXISTENCE, REFERENCE, IMPLICIT, or DESIGNATION, as given by the category name. The default is COMPOSITION.
- Cardinality. As defined by the cardinality range. A missing lower bound is equivalent to a lower bound of 0. A missing upper bound is equivalent to an implementation-defined upper bound (which may depend on the link type) (see clause 24). A missing cardinality is equivalent to [0 ..] if there is a key list and [0..1] if not.
- Exclusiveness: EXCLUSIVE if the keyword '**exclusive**' is present, otherwise SHARABLE.
- Stability: ATOMIC\_STABLE if the stability name '**atomic stable**' is present, COMPOSITE\_STABLE if the stability name '**composite stable**' is present, and NON\_STABLE if no stability name is present.
- Duplication: NON\_DUPLICATED if the keyword '**non\_duplicated**' is present, otherwise DUPLICATED.
- The reverse link type is the link type denoted by the link type name after '**reverse**'; if this is absent, there is no reverse link type if the link type is a designation link type, and otherwise the reverse link type is an implicit link type of cardinality many with no local name.

The new link type in SDS has the following properties (see 8.4.3).

- Destination object types. This is the set of object types denoted by the object type list.
- Key attribute types. This is the sequence of attributes defined by that attribute indication list (see below).
- Non-key attribute types. This is the set of attributes defined by the attribute indication list following the keyword '**with**', or the empty set if there is no such attribute indication list present.
- The usage and export modes are set by the type mode declaration, if present (see clause B.6); the default is usage mode and export mode both set to CREATE, DELETE, and NAVIGATE.
- The maximum usage mode is set to CREATE, DELETE, and NAVIGATE.

When used in a key list, an attribute indication list denotes the list of attribute types as defined above, in the order of the appearance of their attribute type names or attribute type declarations in the attribute indication list.

A *link type extension* extends the link type definition with the same local name in the current SDS section, by adding further outgoing destination object types and attribute types, defined as for a link type declaration.

## B.5 Enumeration types

### Syntax

```
enumeration type declaration =  
    local name, ':', enumeration image, '{', enumeration image};
```

### Meaning

An enumeration type declaration defines one enumerational type, and one enumerational type in SDS in the current SDS, for each of the enumeration images. The local name may be used in an enumeration type name (see clause B.7) in an attribute type declaration to denote the enumeration value type consisting of the defined sequence of enumerational types.

Each enumerational type in SDS has the corresponding enumeration image as its image (see 8.4.4).

## B.6 Type mode declarations

### Syntax

```
type mode declaration = '(', 'usage', type mode, ',', 'export', type mode, ')'  
    | '(', ['usage', ',', 'export'], type mode, ')';  
type mode = 'protected' | allowed access, '{', allowed access};  
allowed access = 'read' | 'write' | 'navigate' | 'create' | 'delete';
```

### Meaning

A type mode declaration occurs in a type declaration, and sets the usage and export modes of the type in SDS denoted by the type declaration to the values denoted by the type modes after **usage** and **export** respectively. The second form of type mode declaration sets both usage and export modes to the values denoted by the type mode.

The type mode denotes the set of definition mode values corresponding to the allowed access values mentioned (i.e. READ\_MODE, WRITE\_MODE, NAVIGATE\_MODE, CREATE\_MODE, and DELETE\_MODE, corresponding to **read**, **write**, **navigate**, **create**, and **delete** respectively), or the empty set if the type mode is **protected**.

## B.7 Names

### Syntax

```
object type name = global name | local name;  
object type list = object type name, '{', object type name};  
attribute type name = global name | local name;  
attribute type list = attribute type name, '{', attribute type name};  
link type name = global name | local name;
```



link type list = link type name, {'', link type name};  
 enumeration type name = global name | local name;  
 sds name = name;  
 local name = name;  
 global name = sds name, '-', local name;  
 name = identifier | ' ' ', character, { character }, ' ' ';  
 identifier = letter, {letter | digit | '\_'};  
 character = (\* any value of the PCTE datatype Character (see 23.1.1.6) \*);  
 newline = (\* implementation-defined \*);  
 capital letter = 'A' | 'B' | 'C' | 'D' | 'E' | 'F' | 'G' | 'H' | 'I' | 'J' | 'K' | 'L' | 'M' | 'N' | 'O' | 'P' | 'Q' | 'R' | 'S' |  
 'T' | 'U' | 'V' | 'W' | 'X' | 'Y' | 'Z';  
 small letter = 'a' | 'b' | 'c' | 'd' | 'e' | 'f' | 'g' | 'h' | 'i' | 'j' | 'k' | 'l' | 'm' | 'n' | 'o' | 'p' | 'q' | 'r' | 's' | 't' | 'u' | 'v' |  
 'w' | 'x' | 'y' | 'z';  
 letter = capital letter | small letter;  
 digit = '0' | '1' | '2' | '3' | '4' | '5' | '6' | '7' | '8' | '9';

### Constraints

An object, attribute, link, or enumeration type name without an SDS name must occur in an object, attribute, link, or enumeration type declaration, respectively, within the local specification.

The local name of an object, attribute, link, or enumeration type name with an SDS name must occur in an object, attribute, link, or enumeration type declaration, respectively, within the specification with that SDS name.

No spaces or control characters may occur between the component parts of a global name, nor within compound lexical items, i.e. keywords and compound symbols ('..', '--', ':=', '""'), enumeration images, initial values (except in strings, where they are significant), lower and upper bounds, and identifiers. One or more separators (spaces and/or newlines) must occur between adjacent keywords and identifiers.

The sequence of characters of a name of the second form must respect the syntax of names in 23.1.3.2.

### Meaning

*Local names* denote object types, attribute types, and link types within an SDS, as established by the corresponding clauses.

*Sds names* denote SDSs; an SDS name may be prefixed to a local name declared in that SDS to disambiguate it.

NOTE – Differences of font and style are not significant; in particular there is no mandatory convention for distinguishing keywords from identifiers. Upper and lower case letters are however distinguished, and underscores in identifiers and keywords are significant, so that 'type1', 'Type1', and 'Type\_1' are all distinct.

## B.8 Comments

### Syntax

comment = '—', { character }, newline;

### Constraints

A comment may contain any graphic characters. It may appear anywhere that a control character can, and nowhere else. It is terminated by the first newline after the opening '--'.

### Meaning

A comment has no effect on the meaning of the DDL definition.

## B.9 Use of DDL identifiers as technical terms

There are two styles for referring to DDL identifiers in running text:

- Style 1. The description is phrased in representational terms, i.e. in terms of objects, attributes, and links. The appropriate style is as in:
  - . A "type\_in\_sds" object is created.
  - . The destination of the "adopted\_user\_group" link from the "system-process" object *process* is ...
  - . ... the value of the "contents\_confidentiality\_label" attribute of the "volume\_object" object must be ...

The SDS name is not given; all DDL object type identifiers are distinct.

- Style 2. The description is phrased in logical terms, i.e. in terms of logical entities. The appropriate style is as in:
  - . A type in SDS is created.
  - . The adopted user group of the process *process* is ...
  - . ... the contents confidentiality label of the device must be ...

These phrases are equivalent. Style 1 is always possible and is used when the mapping between logical entities and objects is not straightforward (e.g. types), at least to describe the mapping. Style 2 is preferred when there is no ambiguity; though the names used, especially for links, sometimes make this style grammatically impossible.

The equivalent phrases 'a "type\_in\_sds" object' and 'a type in SDS' signify an object of type "type\_in\_sds" or of any descendant of "type\_in\_sds"; similarly for other object types.

## Annex C (normative)

### Specification of Errors

#### C.1 Error conditions and error names

For each error condition, the definition is obtained by textually substituting the actual parameters for the corresponding formal parameters throughout the text given in this annex. A list of two or more items as an actual parameter is enclosed in parentheses '( )' with the items separated by commas. A formal parameter in square brackets '[ ]' means that the corresponding actual parameter may be omitted.

Each error name is in two parts: the name of the troublesome entity followed by an indication of what is wrong.

#### C.2 Scope of error conditions

Some error condition definitions take as parameters an object designator *object* and an object scope *scope*. These are used to define a set of objects to any of which the error condition may apply. This set of objects is the *scope* of the error condition and is defined as follows.

- *scope* = ATOMIC. The scope is the object *object*.
- *scope* = COMPOSITE. The scope is the object *object* and all its subcomponents.
- *scope* = COMPONENTS. The scope is an implementation-dependent subset of the object *object* and its components.

#### C.3 Groupings of errors

##### C.3.1 Access errors

Certain error conditions apply to many operations which access an object or objects; they are grouped together for convenience under the name of *access errors*:

ACCESS\_ERRORS (*object*, *scope*, *access\_mode*, [ *permission* ])

If several objects are mentioned, the error conditions apply to each separately. The error conditions represented by ACCESS\_ERRORS depend on the value of *access\_mode*, as follows.

- For *access\_mode* = READ, MODIFY, CHANGE, or SYSTEM\_ACCESS:
  - OBJECT\_IS\_ARCHIVED (*object*)
  - OBJECT\_IS\_INACCESSIBLE (*object*, *scope*)
  - if object resides in a cluster then:
    - ACCESS\_ERRORS (cluster of *object*, *scope*, *access\_mode*, *permission*)
- Additionally, for *access\_mode* = READ, MODIFY, or CHANGE:
  - CONFIDENTIALITY\_WOULD\_BE\_VIOLATED (*object*, *scope*)
  - DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED (*object*, *scope*, [ *permission* ])

- INTEGRITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED (*object, scope*)
- Additionally, for *access\_mode* = MODIFY, CHANGE, or SYSTEM\_ACCESS:
  - VOLUME\_IS\_FULL (volume on which *object* resides)
- Additionally, for *access\_mode* = MODIFY or CHANGE:
  - CONFIDENTIALITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED (*object, scope*)
  - INTEGRITY\_WOULD\_BE\_VIOLATED (*object, scope*)
  - PRIVILEGE\_IS\_NOT\_GRANTED (PCTE\_REPLICATION)
  - REPLICATED\_COPY\_UPDATE\_IS\_FORBIDDEN (*object, scope*)
  - VOLUME\_IS\_READ\_ONLY (volume on which *object* resides)
- Additionally, for *access\_mode* = MODIFY:
  - OBJECT\_IS\_STABLE (*object*)

### C.3.2 Value limit errors

Error conditions relating to the limits on attribute values are grouped together as *value limit errors*:

VALUE\_LIMIT\_ERRORS (*value*)

This is equivalent to one or more cases of LIMIT\_WOULD\_BE\_EXCEEDED, according to the value type of *value*:

- value type is INTEGER:
  - LIMIT\_WOULD\_BE\_EXCEEDED ((MAX\_INTEGER\_ATTRIBUTE, MIN\_INTEGER\_ATTRIBUTE))
- value type is FLOAT:
  - LIMIT\_WOULD\_BE\_EXCEEDED ((MAX\_FLOAT\_ATTRIBUTE, MIN\_FLOAT\_ATTRIBUTE, SMALLEST\_FLOAT\_ATTRIBUTE))
- value type is STRING:
  - LIMIT\_WOULD\_BE\_EXCEEDED (MAX\_STRING\_ATTRIBUTE\_SIZE)
- value type is TIME:
  - LIMIT\_WOULD\_BE\_EXCEEDED ((MAX\_TIME\_ATTRIBUTE, MIN\_TIME\_ATTRIBUTE))

### C.3.3 OWNER right errors

OWNER\_PROPAGATION\_ERRORS\_ON\_COMPONENT\_CREATION (*object*)

This stands for the following, where *group* is the group for which the OWNER mode is being changed:

ATOMIC\_ACL\_IS\_INCOMPATIBLE\_WITH\_OWNER\_CHANGE (*object*)  
 OBJECT\_HAS\_GROUP\_WHICH\_IS\_ALREADY\_OWNER (*object, group*)

COMPONENT\_ADDITION\_ERRORS (*dest*, *link*)

This stands for:

ACCESS\_ERRORS (*dest*, COMPOSITE, CHANGE, OWNER)

ACCESS\_ERRORS (*dest*, ATOMIC, CHANGE, CONTROL\_DISCRETIONARY)

ACCESS\_ERRORS (component of *dest*, ATOMIC, CHANGE, CONTROL\_DISCRETIONARY)

LINK\_EXCLUSIVENESS\_WOULD\_BE\_VIOLATED (*dest*, *link*)

If origin of *link* has OWNER granted or denied for a group *group*:

OBJECT\_HAS\_GROUP\_WHICH\_IS\_ALREADY\_OWNER (*dest*, *group*)

OBJECT\_OWNER\_CONSTRAINT\_WOULD\_BE\_VIOLATED (*dest*)

The following implementation-dependent error may be raised:

OBJECT\_IS\_INACCESSIBLE (outer object of *dest*, ATOMIC)

#### C.4 Other errors

ACCESS\_MODE\_IS\_INCOMPATIBLE (*scope*, *modes*) An attempt is being made to check for OWNER permission in *modes* when the scope *scope* is ATOMIC, or for CONTROL\_DISCRETIONARY permission in *modes* when the scope *scope* is COMPOSITE.

ACCESS\_MODE\_IS\_NOT\_ALLOWED (*scope*, *modes*) The access mode *modes* contains OWNER discretionary access mode and *scope* is ATOMIC or *modes* contains CONTROL\_DISCRETIONARY and *scope* is COMPOSITE.

ACCOUNTING\_LOG\_IS\_NOT\_ACTIVE (*accounting\_log*) *accounting\_log* is not currently an active accounting log.

ACTIVITY\_IS\_OPERATING\_ON\_A\_RESOURCE A resource is currently being operated on by the current activity.

ACTIVITY\_STATUS\_IS\_INVALID (*activity*, *status*) The activity status of *activity* does not have the value *status*.

ACTIVITY\_WAS\_NOT\_STARTED\_BY\_CALLING\_PROCESS The current activity was not started by the calling process.

ARCHIVE\_EXISTS (*archive\_identifier*) The natural value *archive\_identifier* is already the archive identifier of a known archive.

ARCHIVE\_HAS\_ARCHIVED\_OBJECTS (*archive*) The archive *archive* already has archived objects.

ARCHIVE\_IS\_INVALID\_ON\_DEVICE (*device*, *archive*) The contents of *device* either does not correspond in format to the implementation-defined format of an archive, or does not contain the archive identifier of *archive*.

ARCHIVE\_IS\_UNKNOWN (*archive*) The archive *archive* is unknown, i.e. not the destination of a "known\_archive" link from the archive directory.

ATOMIC\_ACL\_IS\_INCOMPATIBLE\_WITH\_OWNER\_CHANGE (*object*) The groups for which OWNER is granted or denied for the origin object of a newly created composition link do

not have CONTROL\_DISCRETIONARY mode granted or denied respectively in the atomic ACL of the destination object *object* of that link.

ATTRIBUTE\_TYPE\_IS\_NOT\_VISIBLE (*reference*) The attribute type identified by *reference* is not a visible type.

ATTRIBUTE\_TYPE\_OF\_LINK\_TYPE\_IS\_NOT\_APPLIED (*reference*) The attribute type of the attribute reference *reference* is not applied to the type of the link relative to which *reference* is being evaluated.

ATTRIBUTE\_TYPE\_OF\_OBJECT\_TYPE\_IS\_NOT\_APPLIED (*reference*) The attribute type of the attribute reference *reference* is not applied to the type of the object relative to which *reference* is being evaluated, nor to any of its visible ancestor types.

AUDIT\_FILE\_IS\_NOT\_ACTIVE (*audit\_file*) The object *audit\_file* is not currently active as the audit file of any station.

BREAKPOINT\_IS\_NOT\_DEFINED (*breakpoint*) *breakpoint* is not a breakpoint set by PROCESS\_ADD\_BREAKPOINT

CARDINALITY\_IS\_INVALID (*link\_type*) The cardinality of *link\_type* is not valid for the operation.

CATEGORY\_IS\_BAD (*object*, *link*, *categories*) The category of the link *link* of the object *object* is not one of *categories*.

CLASS\_NAME\_IS\_INVALID (*name*) The string *name* is not a valid confidentiality or integrity class name (according to context).

CLUSTER\_EXISTS (*cluster\_identifier*, *volume*) The specified cluster number *cluster\_identifier* corresponds to an existing cluster in the volume *volume*.

CLUSTER\_HAS\_OTHER\_LINKS (*cluster*) There are links starting from the cluster *cluster* which are not the "cluster\_in\_volume" link to its associated volume.

CLUSTER\_IS\_UNKNOWN (*cluster*) The "cluster" object *cluster* is not linked to a volume via link of type "known\_cluster".

CONFIDENTIALITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED (*object*, *scope*) A confidentiality confinement violation would occur on one or more objects of the scope defined by *object* and *scope*. More precisely, for some object A of the scope of the error condition:

LABEL\_DOMINATES (confidentiality label of A, confidentiality context of calling process) = **false**

CONFIDENTIALITY\_CRITERION\_IS\_NOT\_SELECTED (*criterion*) The event to be removed in *criterion* is not in the list of confidentiality label dependent criteria.

CONFIDENTIALITY\_LABEL\_IS\_INVALID (*label*) The confidentiality label *label* does not conform to the defined syntax or contains a class that is not a known confidentiality class.

CONFIDENTIALITY\_WOULD\_BE\_VIOLATED (*object*, *scope*) A confidentiality violation would occur on one or more objects of the scope defined by *object* and *scope*. More precisely, for some object A of the scope of the error condition:

LABEL\_DOMINATES (confidentiality context of calling process, confidentiality label of A) = **false**

CONNECTION\_IS\_DENIED The requested connection to the network cannot be made.

CONSUMER\_GROUP\_IS\_IN\_USE (*group*) The consumer group *group* is currently the associated consumer group of a process.

CONSUMER\_GROUP\_IS\_KNOWN (*group*) *group* is already a known consumer group of the accounting directory.

CONSUMER\_GROUP\_IS\_UNKNOWN (*object*) *object* is not a known consumer group of the accounting directory.

CONTENTS\_FORMAT\_IS\_INVALID (*file*) The contents of the file *file* was not saved by QUEUE\_SAVE.

CONTENTS\_IS\_NOT\_EMPTY (*object*) *object* has a non-empty contents.

CONTENTS\_IS\_NOT\_FILE\_CONTENTS (*contents*) The contents handle *contents* does not refer to a file contents.

CONTENTS\_IS\_NOT\_OPEN (*contents*) The contents handle *contents* does not refer to an open object.

CONTENTS\_OPERATION\_IS\_INVALID (*contents\_handle*) *contents\_handle* is the result of an opening where the opening mode or positioning does not allow the current operation to be performed.

CONTROL\_WOULD\_NOT\_BE\_GRANTED (*new\_object*) The CONTROL\_DISCRETIONARY access right or CONTROL\_MANDATORY access right would be no longer granted to any group or would be denied to the predefined user group ALL\_USERS in the atomic ACL of *new\_object* or one of its components.

DATA\_ARE\_NOT\_AVAILABLE (*contents*) It is not possible to read any octets from the opened object *contents*.

DEFAULT\_ACL\_WOULD\_BE\_INCONSISTENT\_WITH\_DEFAULT\_OBJECT\_OWNER (*group*) *group* must have its CONTROL\_DISCRETIONARY access mode value granted in the default atomic ACL of a given process if *group* is the default object owner of the same process.

DEFAULT\_ACL\_WOULD\_BE\_INVALID (*process, group, modes*) No group would have atomic CONTROL\_DISCRETIONARY discretionary access right or atomic CONTROL\_MANDATORY discretionary access right to any objects (as defined in 19.1.3) created by *process* if *group* were set to *modes*.

DEFINITION\_MODE\_VALUE\_WOULD\_BE\_INVALID (*definition\_mode, type*) The definition mode *definition\_mode* contains a mode value which is invalid for the type kind of the type *type* as defined in 8.4.1, 8.4.2, and 8.4.3.

DESTINATION\_OBJECT\_TYPE\_IS\_INVALID (*origin, link, destination\_object*) In the current working schema the object type of *destination\_object* is not among the destination object types of the link type of the link designator *link* of the object *origin*.

DEVICE\_CHARACTERISTICS\_ARE\_INVALID (*device\_characteristics*) The string *device-characteristics* is not a valid device characteristics value.

DEVICE\_CONTROL\_OPERATION\_IS\_INVALID (*device, operation*) The operation *operation* is not a valid operation or the device *device*.

DEVICE\_EXISTS (*device\_identifier*) *device\_identifier* is already the device identifier of a known device.



DEVICE\_IS\_BUSY (*device\_object*, *volume\_identifier*) Another volume than the volume designated by the identifier *volume\_identifier* is mounted on the device represented by *device\_object*.

DEVICE\_IS\_IN\_USE (*device*) The device *device* is in use, i.e. a volume is mounted on it, or its contents is open, or it is in use by an archiving operation.

DEVICE\_IS\_UNKNOWN (*device\_object*) The physical device described by the object *device\_object* does not exist (i.e. there is no "controlled\_device" link from a workstation to *device\_object*).

DEVICE\_LIMIT\_WOULD\_BE\_EXCEEDED (*data*, *contents*) The writing of *data* to *contents* would exceed a device-dependent maximum size.

DEVICE\_SPACE\_IS\_FULL (*device*) The device *device* has insufficient available space.

DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED (*object*, *scope*, [*permission*]) The calling process does not have the atomic or composite right (according to *scope*) *permission* (or one of the rights if a list is given) to the object *object*. If *permission* is absent, the calling process does not have any atomic or composite right (according to *scope*) to *objects*. (If a list of access rights is given, they are all required.)

DISCRETIONARY\_ACCESS\_IS\_NOT\_GRANTED\_TO\_PROCESS (*process*, *object*, *scope*, [*permission* ]) The process *process* does not have the atomic or composite right (according to *scope*) *permission* (or one of the rights if a list is given) to the object *object*. If *permission* is absent, *process* does not have any atomic or composite rights (according to *scope*) to *object*. (If a list of access rights is given, they are all required.)

ENUMERATION\_ATTRIBUTE\_WOULD\_HAVE\_NO\_ENUMERAL\_TYPES (*values*) An enumeration attribute type would be created with an empty sequence *values* of enumerals types.

ENUMERAL\_TYPE\_IS\_INVALID (*value*) *value* is not an enumerals type which is a member of the sequence of enumerals types defined for the enumeration attribute type.

ENUMERAL\_TYPE\_IS\_NOT\_IN\_ATTRIBUTE\_VALUE\_TYPE (*type1*, *type2*) The enumerals type *type1* is not a member of the value type of the enumeration attribute type *type2*.

ENUMERAL\_TYPE\_IS\_NOT\_VISIBLE (*reference*) The enumerals type identified by *reference* is not a visible type.

ENUMERAL\_TYPES\_ARE\_MULTIPLE (*enumerals\_types*) An enumerals type occurs more than once in the sequence *enumerals\_types*.

ENUMERATION\_VALUE\_IS\_OUT\_OF\_RANGE (*initial\_value*, *values*) The value given by *initial\_value* is outside the range of positions defined by the sequence of enumerals *values*.

EVALUATION\_STATUS\_IS\_INCONSISTENT\_WITH\_EVALUATION\_POINT (*reference*, *point*) The evaluation status of the object *reference*, link *reference*, or type *reference* is internal but the evaluation point *point* is not NOW.

EVENT\_TYPE\_IS\_NOT\_SELECTED (*event\_type*) There is no general selection criterion with the selectable event type *event\_type* in the list of criteria.

EXECUTION\_CLASS\_HAS\_NO\_USABLE\_EXECUTION\_SITES (*execution\_class*) The execution class *execution\_class* has no usable execution sites.

EXECUTION\_SITE\_IS\_INACCESSIBLE (*site*) The execution site *site* is not accessible.

EXECUTION\_SITE\_IS\_NOT\_IN\_EXECUTION\_CLASS (*site*, *static\_context*) The execution site *site* is not in the execution class of the static context *static\_context*.

EXECUTION\_SITE\_IS\_UNKNOWN (*site*) The execution site *site* is unknown to the PCTE installation (i.e. it is not linked to the execution site directory with a "known\_execution\_site" link).

EXTERNAL\_LINK\_IS\_BAD (*version*) One or more of the outgoing external existence, reference, or designation links of *version* which is to be copied is to an inaccessible object.

EXTERNAL\_LINK\_IS\_NOT\_DUPLICABLE (*object*) One or more of the outgoing external existence, reference, or designation links of *object* is reversed by a link which is not an implicit link of cardinality many.

FOREIGN\_DEVICE\_IS\_INVALID (*foreign\_device*) *foreign\_device* cannot be interpreted as specifying an appropriate device.

FOREIGN\_EXECUTION\_IMAGE\_HAS\_NO\_SITE (*image*) There is no "on\_foreign\_system" link from the foreign execution image *image*.

FOREIGN\_EXECUTION\_IMAGE\_IS\_BEING\_EXECUTED (*foreign\_system*, *foreign\_name*) The file designated by *foreign\_name* on *foreign\_system* is a foreign execution image which is being executed.

FOREIGN\_OBJECT\_IS\_INACCESSIBLE (*foreign\_system*, *foreign\_name*) The object *foreign\_object* residing on the foreign system *foreign\_system* is not accessible.

FOREIGN\_SYSTEM\_IS\_INACCESSIBLE (*foreign\_system*) The foreign system identified by *foreign\_system* is valid but is not accessible.

FOREIGN\_SYSTEM\_IS\_INVALID (*foreign\_system*, *process*, *class*) The foreign system *foreign\_system*, that is or would be the execution site of *process*, does not support the functionality of the class *class*.

FOREIGN\_SYSTEM\_IS\_UNKNOWN (*foreign\_system*) The foreign system *foreign\_system* does not identify an execution site known to the PCTE installation.

GROUP\_IDENTIFIER\_IS\_IN\_USE (*natural*) The natural value *natural* is already in use as a security group identifier.

GROUP\_IDENTIFIER\_IS\_INVALID (*natural*) The natural value *natural* has never been returned by GROUP\_INITIALIZE.

IMAGE\_IS\_ALREADY\_ASSOCIATED (*image*, *sds*, *enumeration\_type*) The string *image* is already the associated image of an enumeration type in SDS in the SDS *sds*, and there is an enumeration attribute type in SDS in *sds* with value type containing both that enumeration type and *enumeration\_type*.

IMAGE\_IS\_DUPLICATED (*enumeration\_types*, *sds*) Two or more of the enumeration types in SDS associated with the elements of *enumeration\_types* in the SDS *sds* have the same image.

INTEGRITY\_CONFINEMENT\_WOULD\_BE\_VIOLATED (*object*, *scope*) An integrity confinement violation would occur on one or more objects of the scope defined by *object* and *scope*. More precisely, for some object A of the scope of the error condition:

LABEL\_DOMINATES (integrity label of A, integrity context of calling process) = **false**

INTEGRITY\_CRITERION\_IS\_NOT\_SELECTED (*criterion*) The event to be removed in *criterion* is not in the list of integrity label dependent criteria.

INTEGRITY\_LABEL\_IS\_INVALID (*label*) The integrity label *label* does not conform to the defined syntax or contains a class that is not a known integrity class.

INTEGRITY\_WOULD\_BE\_VIOLATED (*object*, *scope*) An integrity violation would occur on one or more objects of the scope defined by *object* and *scope*. More precisely, for some object A of the scope of the error condition:

LABEL\_DOMINATES (integrity context of calling process, integrity label of A) = **false**

INTERPRETER\_IS\_INTERPRETABLE (*static\_context*) The selected interpreter for the static context *static\_context* is itself an interpretable static context and so cannot be used as an interpreter.

INTERPRETER\_IS\_NOT\_AVAILABLE (*static\_context*) *static\_context* is an interpretable static context but no interpreter for it can be executed.

KEY\_ATTRIBUTE\_TYPE\_APPLY\_IS\_FORBIDDEN (*attribute\_type*) The attribute type *attribute\_type* is a key attribute type and cannot be applied.

KEY\_IS\_BAD (*object*, *reference*) The link reference *reference* has a key which supplies either too many or too few key attribute values or supplies key attribute values of the wrong value type for the link type specified by reference and the origin *object*.

KEY\_SYNTAX\_IS\_WRONG (*key*) The text value *key* does not have the syntax of a key.

KEY\_TYPE\_IS\_BAD (*type*) The value type of the attribute type *type* is not natural or string.

KEY\_TYPES\_ARE\_MULTIPLE (*types*) An attribute type occurs more than once in *types*.

KEY\_UPDATE\_IS\_FORBIDDEN (*object*, *link*, *attribute*) The attribute *attribute* of the link *link* of the object *object* is a key attribute and cannot be modified.

KEY\_VALUE\_DOES\_NOT\_EXIST (*link\_reference*, *index*) The index *index* designates a key value in the key of the link reference *link\_reference* which does not exist.

LABEL\_IS\_OUTSIDE\_RANGE (*object*, *device*) A mandatory label of the object *object* is or would be outside the corresponding mandatory security range of the multi-level secure device *device*. More precisely, one of the following is **false**:

CONFIDENTIALITY\_LABEL\_WITHIN\_RANGE (*object*, *device*)

INTEGRITY\_LABEL\_WITHIN\_RANGE (*object*, *device*)

LABEL\_RANGE\_IS\_BAD (*high\_label*, *low\_label*) *high\_label* does not dominate *low\_label* in a new security range for a multi-level secure device. More precisely, LABEL\_DOMINATES (*high\_label*, *low\_label*) is **false**.

LAN\_ERROR\_EXISTS An error has occurred on the local area network.

LIMIT\_WOULD\_BE\_EXCEEDED (*limit*) The implementation limit *limit* would be exceeded.

LINK\_DESTINATION\_DOES\_NOT\_EXIST (*link*) The destination of the link *link* does not exist and so cannot be accessed.

LINK\_DESTINATION\_IS\_NOT\_VISIBLE (*link*) The object type of the destination of the link *link* is not a visible destination object type of the link type of *link*.

LINK\_DOES\_NOT\_EXIST (*object*, *link\_name*) The link specified by *link\_name* and *object* as origin does not exist and therefore cannot be accessed.

LINK\_EXCLUSIVENESS\_WOULD\_BE\_VIOLATED (*dest*, *link*) Creation of the link *link* to the object *dest* would violate exclusivity: either *link* is an exclusive composition link and there is already a composition link to *dest*, or *link* is a composition link and there is already an exclusive composition link to *dest*.

LINK\_EXISTS (*object*, *link*) The link *link* of the object *object* already exists and therefore cannot be created.

LINK\_NAME\_IS\_TOO\_LONG\_IN\_CURRENT\_WORKING\_SCHEMA (*link\_name*) The link name *link\_name* cannot be represented in the current working schema because the limit MAX\_LINK\_NAME\_SIZE would be exceeded.

LINK\_NAME\_SYNTAX\_IS\_WRONG (*link\_name*) The syntax of the link name *link\_name* is wrong.

LINK\_REFERENCE\_IS\_UNSET (*reference*) The link reference *reference* has never been set or has been explicitly unset by LINK\_REFERENCE\_UNSET.

LINK\_TYPE\_CATEGORY\_IS\_BAD (*link\_type*, *categories*) The category of the link type *link\_type* is not one of *categories*.

LINK\_TYPE\_IS\_NOT\_APPLIED\_TO\_OBJECT\_TYPE (*object\_reference*, *link\_reference*) The link type of *link\_reference* is not a visible type of the object type in working schema of the object *object\_reference*.

LINK\_TYPE\_IS\_NOT\_VISIBLE (*reference*) The link type identified by *reference* is not a visible type.

LINK\_TYPE\_IS\_UNKNOWN (*name*) *name* is not the name of a link type in the current working schema.

LINK\_TYPE\_PROPERTIES\_AND\_KEY\_TYPES ARE\_INCONSISTENT (*link\_type\_properties*, *key\_attribute\_types*) The properties specified by *link\_type\_properties* are inconsistent with the key attribute types specified by *key\_attribute\_types*.

LINK\_TYPE\_PROPERTIES ARE\_INCONSISTENT (*link\_type\_properties*) The properties specified by *link\_type\_properties* are inconsistent among themselves (see 8.3.3).

LOCK\_COULD\_NOT\_BE\_ESTABLISHED (*resource*, *scope*) A lock could not be immediately established or promoted on the resource *resource* with scope *scope*.

LOCK\_INTERNAL\_MODE\_CANNOT\_BE\_CHANGED (*object*, *lock\_mode*) There is a conflict between the requested internal lock mode *lock\_mode* of object *object* and other concurrent acquisitions of resources in the concerned domain of *object*.

LOCK\_IS\_NOT\_EXPLICIT (*object*) The object *object* is not explicitly locked by the current activity.

LOCK\_MODE\_IS\_NOT\_ALLOWED (*lock\_mode*) The lock mode *lock\_mode* is WTR or DTR which is not allowed in the current activity.

LOCK\_MODE\_IS\_TOO\_STRONG (*lock\_mode*, *resource*) The required internal mode *lock\_mode* is stronger than the external one established on *resource*.

LOWER\_BOUND\_WOULD\_BE\_VIOLATED (*object, link*) The object *object* is not deleted by the operation, and the deletion of *link* from its origin *object* would leave the number of links of *object* with the link type of *link* less than the lower bound of the link type of *link*.

MANDATORY\_CLASS\_IS\_ALREADY\_DOMINATED (*object*) There already exists a confidentiality or integrity class which directly dominates the confidentiality or integrity class designated by *object*. Only one such class is permitted.

MANDATORY\_CLASS\_IS\_KNOWN(*object*) *object* is already a known mandatory class.

MANDATORY\_CLASS\_IS\_UNKNOWN (*class*) *class* does not identify a mandatory class known to the PCTE installation, i.e. there is no link from the mandatory directory to *class*.

MANDATORY\_CLASS\_NAME\_IS\_IN\_USE (*class\_name*) There already exists a mandatory class with the name *class\_name*.

MASTER\_IS\_INACCESSIBLE (*object, scope*) The master of the object *object*, or *object* itself if normal, is inaccessible.

MAXIMUM\_USAGE\_MODE\_WOULD\_BE\_EXCEEDED (*type, definition\_mode*) The usage mode or export mode *definition\_mode* exceeds the maximum usage mode associated with the type *type*.

MEMORY\_ADDRESS\_IS\_OUT\_OF\_PROCESS (*address, process*) The address *address* is not a valid address in the memory space of the process *process*.

MEMORY\_REGION\_IS\_NOT\_IN\_PROFILING\_SPACE (*start, end*) The addresses *start* and *end* do not define a region in the profiling space of the calling process.

MESSAGE\_IS\_NOT\_A\_NOTIFICATION\_MESSAGE (*message*) The message type of the message *message* is not a notification message type.

MESSAGE\_POSITION\_IS\_NOT\_VALID (*context, queue*) The natural *position* does not denote a position in the message queue *queue*.

MESSAGE\_QUEUE\_HAS\_BEEN\_DELETED (*queue*) The message queue *queue* was deleted while the calling process was waiting to receive a message from it.

MESSAGE\_QUEUE\_HAS\_BEEN\_WOKEN (*queue*) A message of message type WAKE was received while the calling process was waiting to receive a message from the message queue *queue*.

MESSAGE\_QUEUE\_HAS\_NO\_HANDLER (*queue*) No valid handler routine has been specified for the message queue *queue*.

MESSAGE\_QUEUE\_IS\_BUSY (*queue*) There are messages on the message queue *queue*.

MESSAGE\_QUEUE\_IS\_NOT\_RESERVED (*queue*) The caller has not reserved the message queue *queue*.

MESSAGE\_QUEUE\_IS\_RESERVED (*queue*) The message queue *queue* is reserved by another process.

MESSAGE\_QUEUE\_TOTAL\_SPACE\_WOULD\_BE\_TOO\_SMALL (*queue, total\_space*) The space currently in use by message queue *queue* is greater than the requested space limit of *total\_space* or *total\_space* is less than four times MAX\_MESSAGE\_SIZE.

MESSAGE\_QUEUE\_WOULD\_BE\_TOO\_BIG (*queue, file*) The size of the message queue *queue* would exceed the queue's current value of the total space.

MESSAGE\_TYPES\_NOT\_FOUND\_IN\_QUEUE (*queue, types, context*) There is no message of a message type in *types* in the message queue *queue* after the position given by *context*.

NON\_BLOCKING\_IO\_IS\_INVALID (*object, non\_blocking\_io*) An attempt is being made to open an object *object* of type "file" with *non\_blocking\_io* **false**, or to open an object *object* of type "pipe" or "device" with *non\_blocking\_io* **true** when it does not support non-blocking input-output.

NOTIFIER\_KEY\_DOES\_NOT\_EXIST (*natural*) No notifier with the key *natural* exists.

NOTIFIER\_KEY\_EXISTS (*natural*) A notifier with the key *natural* exists.

NUMBER\_OF\_PARAMETERS\_IS\_WRONG (*operation*) The number of parameters of the operation *operation* does not match the operation signature.

OBJECT\_ARCHIVING\_IS\_INVALID (*objects*) One or more of the objects of *objects* cannot be archived, i.e. it is one of the following:

- a master or copy object;
- a volume;
- a locked object;
- a message queue that is reserved or contains one or more messages;
- a pipe or device;
- an active audit file or accounting log;
- a static context that is being executed or interpreted;
- an active process or activity;
- a mandatory class;
- an archive.

OBJECT\_CANNOT\_BE\_CLUSTERED (*object*) An attempt is being made to create an object *object* or a copy of an object *object* in a cluster, or to move an object *object* into a cluster, but the type of *object* is "file", "pipe", "message\_queue", "device", "accounting\_log", "audit\_file", "volume", "cluster", "archive", "archive\_directory", "process", "activity", "common\_root", "sds", "workstation", "execution\_class", "execution\_site", "execution\_site\_directory", "replica\_set\_directory", "replica\_set", "security\_group", "program\_group", "mandatory\_directory", "mandatory\_class", or "security\_group\_directory", or a descendant of one of those types.

OBJECT\_CANNOT\_BE\_STABILIZED (*object*) The object *object* cannot be stabilized, i.e. it is one of the following:

- an active process or activity;
- an active audit file or accounting log;
- a mounted volume;
- a message queue;



- a pipe.

OBJECT\_CRITERION\_IS\_NOT\_SELECTED (*criterion*) The event to be removed in *criterion* is not in the list of object-dependent criteria.

OBJECT\_HAS\_COPIES (*object*) The master object *object* has copies and therefore cannot be removed.

OBJECT\_HAS\_EXTERNAL\_LINKS\_PREVENTING\_DELETION (*object*) The object *object* or a component of *object* that should be deleted by the operation has incoming external reference links or outgoing external existence links.

OBJECT\_HAS\_GROUP\_WHICH\_IS\_ALREADY\_OWNER (*object*, *group*) An attempt is being made to change the OWNER discretionary access mode value to DENIED for *group* when *group* is not in the effective discretionary groups for the calling process, and *group* has OWNER right granted in the composite ACL of *object*. See 19.1.2.

OBJECT\_HAS\_INTERNAL\_LINKS\_PREVENTING\_DELETION (*object*) The object *object* or a component of *object* cannot be deleted as there are internal reference links from another component which is the destination of an external composition or existence link.

OBJECT\_HAS\_LINKS\_PREVENTING\_DELETION (*object*) Other reference links to the object *object* or composition or existence links from it exist.

OBJECT\_IS\_A\_PROCESS (*object*). The object *object* is a process.

OBJECT\_IS\_A\_REPLICA\_SET (*object*) *object* is a replica set and so must be replicated in itself. A copy of *object* must exist on each of its copy volumes.

OBJECT\_IS\_ALREADY\_IN\_RESOURCE\_GROUP (*object*, *group*) The object *object* is already a member of the resource group *group*.

OBJECT\_IS\_ARCHIVED (*object*) The object *object* resides in an archive.

OBJECT\_IS\_FINE\_GRAIN (*object*) *object* is fine-grain and an attempt is being made to perform one of the operations which are not permitted on fine-grain objects.

OBJECT\_IS\_IN\_USE\_FOR\_DELETE (*object*) The object *object*, or any of its non-shared components, is being operated on or is one of:

- the "process" object of an executing process;
- a static context being executed or interpreted;
- the "activity" object of a non-terminated activity;
- a message queue which is non-empty or reserved.

OBJECT\_IS\_IN\_USE\_FOR\_MOVE (*object*) The object *object* or any of its components is being operated on or is one of:

- the "process" object of an executing process;
- a static context being executed or interpreted;
- the "activity" object of a non-terminated activity;
- a message queue which is non-empty or reserved.
- *object* is locked (including an object with open contents);



- *object* is an active accounting log and accounting is switched on;
- *object* is an active audit file.

OBJECT\_IS\_INACCESSIBLE (*object*, *scope*) One or more of the objects defined by *object* and *scope* is not accessible.

OBJECT\_IS\_INACCESSIBLY\_ARCHIVED (*object*, *scope*) One or more of the objects defined by *object* and *scope* is archived and the archive is inaccessible.

OBJECT\_IS\_LOCKED (*object*, *scope*) One or more objects defined by *object* and *scope* is locked.

OBJECT\_IS\_NOT\_ACCOUNTABLE\_RESOURCE (*object*) The object *object* is not an accountable resource.

OBJECT\_IS\_NOT\_ARCHIVED (*object*) The object *object* does not reside in an archive.

OBJECT\_IS\_NOT\_CONVERTIBLE (*object*) The object *object* cannot have its type converted because its replicated state is not NORMAL, and it is an object of a type such that it cannot be replicated (see 17.1.2).

OBJECT\_IS\_NOT\_IN\_RESOURCE\_GROUP (*object*, *group*) The object *object* is not in the resource group *group*.

OBJECT\_IS\_NOT\_LOCKED (*object*) There is no lock on the object resource *object*.

OBJECT\_IS\_NOT\_MASTER\_REPLICATED\_OBJECT (*object*) The object *object* does not have replication status MASTER.

OBJECT\_IS\_NOT\_MOVABLE (*object*, *scope*) The type of one or more of the objects of the scope defined by *object* and *scope* is one of the predefined types "volume", "device", and "workstation", or a descendant of one of those types.

OBJECT\_IS\_NOT\_ON\_ADMINISTRATION\_VOLUME (*object*, *station*) The object *object* does not reside on the administration volume of the workstation *station*.

OBJECT\_IS\_NOT\_ON\_MASTER\_VOLUME\_OF\_REPLICA\_SET (*replica\_set*, *object*) The object *object* does not reside on the master volume of the replica set *replica\_set*.

OBJECT\_IS\_NOT\_REPLICABLE (*object*) The type of the object *object* is one of the predefined types "process", "activity", "pipe", "message\_queue", "volume", "audit\_file", "accounting\_log", "device", and "execution\_site", or a descendant of one of them.

OBJECT\_IS\_NOT\_REPLICATED\_ON\_VOLUME (*object*, *volume*) The object *object* is not a replicated object with a master or copy replica on the volume *volume*.

OBJECT\_IS\_OF\_WRONG\_TYPE (*reference*) The object type of the object *reference* is not the object type required by the operation, nor a descendant of that type.

OBJECT\_IS\_OPERATED\_ON (*object*, *scope*) One or more objects of the scope defined by *object* and *scope* which is to be unlocked is currently being operated on.

OBJECT\_IS\_PREDEFINED\_REPLICATED (*object*) The object *object* is a predefined replicated object.

OBJECT\_IS\_REPLICATED (*object*) The object *object* is either a master or copy of a replicated object.

OBJECT\_IS\_STABLE (*object*) The object *object* is the destination of an atomically stabilizing link, or is a component of the destination of a compositely stabilizing link.

OBJECT\_LABEL\_CANNOT\_BE\_CHANGED\_IN\_TRANSACTION (*object*) The object *object* is a message queue or a pipe and an attempt is being made to change the confidentiality or integrity label of *object* on behalf of a transaction.

OBJECT\_OWNER\_CONSTRAINT\_WOULD\_BE\_VIOLATED (*object*) An attempt to change the atomic or composite ACL of *object* would result in an inconsistency with the OWNER rights on *object* or an outer object of *object*.

OBJECT\_OWNER\_VALUE\_WOULD\_BE\_INCONSISTENT\_WITH\_ATOMIC\_ACL (*object*) An attempt has been made to set up an atomic ACL for newly created *object* with CONTROL\_DISCRETIONARY mode not granted for a discretionary group which is the default object owner and which therefore has OWNER mode granted in the composite ACL.

OBJECT\_REFERENCE\_IS\_INTERNAL (*reference*) *reference* is an internal object reference.

OBJECT\_REFERENCE\_IS\_INVALID (*reference*) The reference *reference* designates an object which has been moved, archived, or deleted.

OBJECT\_REFERENCE\_IS\_UNSET (*reference*) The object reference *reference* has never been set or has been explicitly unset by REFERENCE\_UNSET.

OBJECT\_TYPE\_IS\_ALREADY\_IN\_DESTINATION\_SET (*link\_type*, *object\_type*, *sds*) The object type *object\_type* is already in the destination object types of *link\_type* in the SDS *sds*.

OBJECT\_TYPE\_IS\_INVALID (*type*) The object type *type* is "volume", "device", or a descendant of "volume" or "device".

OBJECT\_TYPE\_IS\_NOT\_IN\_DESTINATION\_SET (*link\_type*, *object\_type*, *sds*) The object type *object\_type* is not in the destination object types of *link\_type* in the SDS *sds*.

OBJECT\_TYPE\_IS\_NOT\_VISIBLE (*reference*) The object type identified by *reference* is not a visible object type.

OBJECT\_TYPE\_IS\_UNKNOWN (*type*) *type* is not an object type in the current working schema.

OBJECT\_TYPE\_WOULD\_HAVE\_NO\_PARENT\_TYPE (*parents*) An object type would be created with an empty set *parents* of parent types.

OPEN\_KEY\_IS\_INVALID (*key*). The key *key* is not a valid open key because there already exists an "open\_object" link with key *key* from the calling process, or because it is greater than MAX\_OPEN\_OBJECTS\_PER\_PROCESS.

OPENING\_MODE\_IS\_INVALID (*object*, *opening\_mode*) *opening\_mode* is not compatible with the type of *object*

OPERATION\_HAS\_TIMED\_OUT The duration of the operation has exceeded the time out of the calling process at the time the operation was called.

OPERATION\_IS\_INTERRUPTED The current process has been interrupted while executing the current operation.

OPERATION\_IS\_NOT\_ALLOWED\_ON\_TYPE (*reference*) The operation has used a type which is not visible and the operation is one of the operations where its usage is not allowed, even when PCTE\_CONFIGURATION is an effective group of the process.

OPERATION\_METHOD\_CANNOT\_BE\_FOUND (*operation*) The method related to the operation *operation* cannot be found in the method repository.

OPERATION\_METHOD\_CANNOT\_BE\_ACTIVATED (*operation*) The method related to the operation *operation* cannot be activated.

PARENT\_BASIC\_TYPES\_ARE\_MULTIPLE (*parent\_types*) *parent\_types* contains types which are or are descended from at least two different types from the following list: "system-file", "system-pipe", "system-device", "system-volume", "system-message\_queue", "system-process", "system-activity", "security-audit\_file", "accounting-accounting\_log"; or which are or are descended from at least two different types from the following list: "security-user", "security-user\_group", "security-program\_group".

PATHNAME\_SYNTAX\_IS\_WRONG (*pathname*) The syntax of the text value *pathname* is not that of a pathname.

PIPE\_HAS\_NO\_WRITERS (*contents*) An attempt has been made to read the blocking pipe *contents* and no contents handle was open in APPEND\_ONLY mode when the operation was called.

POSITION\_HANDLE\_IS\_INVALID (*position\_handle*, *contents\_handle*) *position\_handle* does not designate a valid position in the sequence of octets designated by *contents\_handle*.

POSITION\_IS\_INVALID (*position*) The position *position* is less than FIRST.

POSITIONING\_IS\_INVALID (*contents*, *positioning*) The required positioning *positioning* is invalid for the specified file or device.

PREFERENCE\_DOES\_NOT\_EXIST (*object*, *reference*) Evaluation of the link reference *reference* relative to the origin object *object* requires some preferred key attributes or some preferred link type to be specified, but the required preference information is not provided by *object*.

PREFERRED\_LINK\_KEY\_IS\_BAD (*object*, *link*, *string*) The string *string* is unsuitable as the preferred key type for the link *link* of the object *object*; it supplies too many or too few key attribute types, supplies key value attributes of the wrong types, or is syntactically invalid for a preferred key.

PREFERRED\_LINK\_TYPE\_IS\_UNSET (*object*) The preferred link type of the object *object* is already unset.

PRIVILEGE\_IS\_NOT\_GRANTED (*group*) The predefined security group *group* is not effective for the calling process.

PROCESS\_CONFIDENTIALITY\_IS\_NOT\_DOMINATED (*confidentiality\_label*, *process*) The confidentiality label of the process *process* is not dominated by *confidentiality\_label*.

PROCESS\_FILE\_SIZE\_LIMIT\_WOULD\_BE\_EXCEEDED (*data*, *contents*) The writing of *data* to the file *contents* would cause *contents* to exceed the process file size limit for the calling process.

PROCESS\_HAS\_NO\_UNTERMINATED\_CHILD The calling process has no child process.

PROCESS\_INTEGRITY\_DOES\_NOT\_DOMINATE (*integrity\_label*, *process*) The integrity label of the process *process* does not dominate *integrity\_label*.

PROCESS\_IS\_IN\_TRANSACTION The calling process is currently running in a transaction or the calling process is currently running in an activity nested in a transaction.

PROCESS\_IS\_INITIAL\_PROCESS (*process*) The process *process* is the initial process of a workstation, and so cannot be terminated.

PROCESS\_IS\_NOT\_ANCESTOR (*process*) The process *process* is neither the calling process nor an ancestor of the calling process.

PROCESS\_IS\_NOT\_CHILD (*process*) The process *process* is not a child of the calling process.

PROCESS\_IS\_NOT\_TERMINABLE\_CHILD (*process*) The process *process* is either in a READY state or is not a child of the current process.

PROCESS\_IS\_NOT\_THE\_CALLER (*process*) The process *process* is not the calling process.

PROCESS\_IS\_THE\_CALLER (*process*) The process *process* is the calling process.

PROCESS\_IS\_UNKNOWN (*process*) The process *process* has process status UNKNOWN.

PROCESS\_LABELS\_WOULD\_BE\_INCOMPATIBLE(*user*) The determination of the new mandatory labels of the calling process has failed because the new label values are incompatible with the clearance of user *user* or the security ranges of the workstation or volume on which the calling process is located (see 13.4.11).

PROCESS\_LACKS\_REQUIRED\_STATUS (*process*, *status*) The process *process* has not got the required status *status* (or any of the list of statuses *status*), other than UNKNOWN.

PROCESS\_TERMINATION\_IS\_ALREADY\_ACKNOWLEDGED (*process*) The acknowledged termination of *process* is already **true**.

PROFILING\_IS\_NOT\_SWITCHED\_ON Profiling is not yet switched on for the calling process.

PROGRAM\_GROUP\_IS\_NOT\_EMPTY (*group*) There are static contexts in the program group *group* so that it cannot be removed.

RANGE\_IS\_OUTSIDE\_RANGE (*device1*, *device2*) The mandatory security range of the device *device1* is or would be not entirely enclosed within the mandatory security range of the multi-level secure device *device2*. More precisely, one of the following does not hold:

CONFIDENTIALITY\_RANGE\_WITHIN\_RANGE (*device1*, *device2*)

INTEGRITY\_RANGE\_WITHIN\_RANGE (*device1*, *device2*)

REFERENCE\_CANNOT\_BE\_ALLOCATED There is not enough space to create an internal reference.

REFERENCE\_NAME\_IS\_INVALID (*reference\_name*) *reference\_name* is not a valid referenced object name.

REFERENCED\_OBJECT\_IS\_NOT\_MUTABLE (*referenced\_object*) The referenced object *referenced\_object* is non-mutable (e.g. \$common\_root, \$static\_context).

REFERENCED\_OBJECT\_IS\_UNSET(*reference*) The object reference *reference* contains a referenced object name which is neither the key of a "referenced\_object" link from the calling process, nor an alias.

RELATIONSHIP\_TYPE\_PROPERTIES\_ARE\_INCONSISTENT (*forward\_properties, backward\_properties*) The forward link type properties *forward\_properties* and the reverse link type properties *backward\_properties* are inconsistent with the properties of a link type and its reverse link type (see 8.3.3).

REPLICA\_SET\_COPY\_IS\_NOT\_EMPTY (*replica\_set, volume*) The volume *volume* contains copies of objects in the replica set *replica\_set* other than a copy of *replica\_set* itself.

REPLICA\_SET\_HAS\_COPY\_VOLUMES (*replica\_set*) The replica set *replica\_set* has at least one associated copy volume.

REPLICA\_SET\_IS\_NOT\_EMPTY (*replica\_set*) The replica set *replica\_set* contains masters of objects other than the master of *replica\_set* itself.

REPLICA\_SET\_IS\_NOT\_KNOWN (*replica\_set*) The replica set *replica\_set* is not known within the replica set directory.

REPLICATED\_COPY\_IS\_IN\_USE (*object*) The copy object *object* has a usage designation link.

REPLICATED\_COPY\_UPDATE\_IS\_FORBIDDEN (*object, scope*) An object of the scope defined by *object* and *scope* is replicated and the calling process is attempting to update a copy.

RESOURCE\_GROUP\_IS\_KNOWN (*object*) The object *object* is already known by the PCTE installation as a valid resource group.

RESOURCE\_GROUP\_IS\_UNKNOWN (*group*) *group* does not identify a known resource group of the accounting directory.

REVERSE\_KEY\_IS\_BAD (*origin, link, destination, reverse\_key*) The link *link* of the object *origin* is reversed by a link of cardinality many for which *reverse\_key* supplies either too many or too few key attribute values or supplies key attribute values of the wrong value type.

REVERSE\_KEY\_IS\_NOT\_SUPPLIED (*origin, link, destination*) The link *link* of the object *origin* has a reverse link of cardinality many for which no key can be inferred, i.e. it is not an implicit link, *destination* does not have a preferred key specified for its type, and no specific key value is provided as a parameter to the operation.

REVERSE\_KEY\_IS\_SUPPLIED (*reverse\_key*) A reverse key *reverse\_key* is supplied for the reverse of a created link though the reverse link is of category IMPLICIT or cardinality one..

REVERSE\_LINK\_EXISTS (*origin, link, destination, reverse\_key*) The link *link* of the object *origin* is reversed by a link of cardinality many for which *reverse\_key* would result in the creation of a link which already exists in the links of *destination*.

SDS\_IS\_IN\_A\_WORKING\_SCHEMA (*sds\_name*) The SDS *sds\_name* is currently included in a working schema (i.e there is an "in\_working\_schema\_of" link from *sds\_name*)

SDS\_IS\_KNOWN (*sds*) The SDS *sds* is already known to the PCTE installation.

SDS\_IS\_NOT\_EMPTY\_NOR\_VERSION (*sds*) The "sds" object *sds* is neither empty (i.e. has no types in SDS) nor a version of a known SDS. An object is a version of another object if related by a series of "predecessor" or a series of "successor" links.

SDS\_IS\_NOT\_IN\_WORKING\_SCHEMA (*sds*) The SDS *sds* is not in the current working schema.



SDS\_IS\_PREDEFINED (*sds*) The SDS *sds* is predefined and cannot be changed.

SDS\_IS\_UNDER\_MODIFICATION (*sds\_name*) The SDS *sds\_name* contains typing information currently being modified in an uncommitted non-enclosing transaction.

SDS\_IS\_UNKNOWN (*sds*) The SDS designator *sds* does not identify a known SDS in the SDS directory.

SDS\_NAME\_IS\_DUPLICATE (*sds\_name*) There is already an SDS of the name *sds\_name*.

SDS\_NAME\_IS\_INVALID (*name*) *name* is not valid for an SDS name.

SDS\_WOULD\_APPEAR\_TWICE\_IN\_WORKING\_SCHEMA (*sds\_sequence*) *sds\_sequence* has two SDS names of the same "sds" object. This would result in an invalid working schema.

SECURITY\_GROUP\_ALREADY\_HAS\_THIS\_SUBGROUP (*group*, *subgroup*) *subgroup* is already a subgroup of *group*.

SECURITY\_GROUP\_IS\_ALREADY\_ENABLED (*group*, *class*) *group* is already enabled for confidentiality downgrade/integrity upgrade from *class*.

SECURITY\_GROUP\_IS\_IN\_USE (*group*) The group *group* is effective for a process.

SECURITY\_GROUP\_IS\_KNOWN (*group*) *group* is already defined as a security group.

SECURITY\_GROUP\_IS\_NOT\_A\_SUBGROUP (*subgroup*, *group*) *subgroup* is not a subgroup of *group*.

SECURITY\_GROUP\_IS\_NOT\_ADOPTABLE (*user\_group*, *process*) *user\_group* is not an adoptable user group for the process *process*.

SECURITY\_GROUP\_IS\_NOT\_ENABLED (*group*, *class*) *group* is not enabled for confidentiality downgrade/integrity upgrade from *class*.

SECURITY\_GROUP\_IS\_PREDEFINED (*group*) *group* is one of the predefined security groups (see 19.1.1).

SECURITY\_GROUP\_IS\_REQUIRED\_BY\_OTHER\_GROUPS (*group*) *group* has subgroups or is a subgroup or (for a user group or program group) has group members or (for a user) is a member of user groups.

SECURITY\_GROUP\_IS\_UNKNOWN (*group*) *group* does not identify a security group known to the PCTE installation i.e. there is no link from the security group directory.

SECURITY\_GROUP\_WOULD\_BE\_IN\_INVALID\_GRAPH (*subgroup*, *group*) *subgroup* may not become a subgroup of *group* since this would result in a graph of security groups which is not a directed acyclic graph.

SECURITY\_POLICY\_WOULD\_BE\_VIOLATED A violation of the built-in security policy of the PCTE implementation (see 20.1.8) has been attempted.

STATIC\_CONTEXT\_CONTENTS\_CANNOT\_BE\_EXECUTED (*static\_context*, *station*) The contents of the static context *static\_context* cannot be executed on the workstation *station*.

STATIC\_CONTEXT\_IS\_ALREADY\_MEMBER (*program*, *group*) The static context *program* is already a member of the program group *group*.

STATIC\_CONTEXT\_IS\_BEING\_WRITTEN (*static\_context*) The object contents of *static\_context* is open for writing.

STATIC\_CONTEXT\_IS\_IN\_USE (*object*) *object* is a static context that is being executed or interpreted.

STATIC\_CONTEXT\_IS\_NOT\_MEMBER (*program*, *program\_group*) *program* is not a member of the program group *program\_group*.

STATIC\_CONTEXT\_REQUIRES\_TOO\_MUCH\_MEMORY (*static\_context*) Execution of *static\_context* requires more memory than is available.

STATUS\_IS\_BAD (*status*) The value of *status* is neither CONNECTED nor CLIENT.

TIME\_CANNOT\_BE\_CHANGED The system time cannot be changed.

TRANSACTION\_CANNOT\_BE\_COMMITTED The activity to be committed is a transaction and the system cannot commit the updates made on behalf of this transaction. Note that this error implies an ABORT\_ACTIVITY. This may be due, for example, to network failure, system crash, disk crash, disk physically switched to a read only protection etc. The circumstances under which this occurs are implementation-defined.

TYPE\_CANNOT\_BE\_APPLIED\_TO\_LINK\_TYPE (*link\_type*, *attribute\_type*). The attribute type *attribute\_type* cannot be applied to the link type *link\_type* because it is the "object\_on\_volume" link type or a usage designation or service designation link type.

TYPE\_HAS\_DEPENDENCIES (*sds*, *type*) The type *type* cannot be deleted because it has type dependencies and its deletion would violate SDS well-formedness rules, i.e. within the SDS *sds*:

- for an attribute type, it is applied to an object or link type;
- for a link type, it is applied to an object type, or it has a non-empty set of destination link types, or non-key attribute types are applied to it; or the same is true of its reverse link type, if any;
- for an object type, it has a child type, or there is an attribute or link type applied to it;
- for an enumerational type, it is associated with an enumeration attribute type.

TYPE\_HAS\_NO\_LOCAL\_NAME (*sds*, *type\_reference*) The type in *sds* identified by *type\_reference* does not have a local name in the SDS *sds*.

TYPE\_IDENTIFIER\_IS\_INVALID (*identifier*) *identifier* is not a valid type identifier.

TYPE\_IDENTIFIER\_USAGE\_IS\_INVALID (*reference*) The operation does not allow a type identifier such as *reference* to be used.

TYPE\_IS\_ALREADY\_APPLIED (*sds*, *type*, *definition\_type*) In the SDS *sds*, the type (either an attribute type or a link type) *type* is already applied to *definition\_type* (either a link type or an object type).

TYPE\_IS\_ALREADY\_CONSTRAINED (*sds*, *parameter\_type*) The parameter type *parameter\_type* is already constrained to an attribute type, an object type, or an interface type.

TYPE\_IS\_ALREADY\_KNOWN\_IN\_SDS (*type*, *to\_sds*) The type *type* is already defined in the SDS *sds*.

TYPE\_IS\_NOT\_APPLIED (*sds*, *type1*, *type2*) The type *type1* is not applied to the type *type2* in the SDS *sds*.



TYPE\_IS\_NOT\_DESCENDANT (*type*, *other\_type*) The object type *other\_type* is neither *type* nor a descendant object type of *type*.

TYPE\_IS\_NOT\_VISIBLE (*type*) The type reference or type name *type* is not visible.

TYPE\_IS\_OF\_WRONG\_KIND (*reference*) The reference *reference* identifies a type which is not of the kind expected by the operation (object type, link type, attribute type, or enumerational type).

TYPE\_IS\_UNKNOWN\_IN\_SDS (*sds*, *type*) The type *type* is not defined in the SDS *sds*.

TYPE\_NAME\_IN\_SDS\_IS\_DUPLICATE (*sds\_name*, *type\_name*) The name *type\_name* duplicates a local name of a type already defined in the SDS *sds\_name*.

TYPE\_NAME\_IS\_INVALID (*type\_name*) *type\_name* does not conform to the syntax of a type name.

TYPE\_OF\_OBJECT\_IS\_INVALID (*object*, *scope*) The type of one or more of the objects of the scope defined by *object* and *scope* is "volume", "device", or a descendant of "volume" or "device".

TYPE\_OF\_PARAMETER\_IS\_WRONG (*operation*, *parameter*) The type of the value of the parameter *parameter* does not match that defined by the signature of the operation *operation*.

TYPE\_REFERENCE\_IS\_INVALID (*reference*) The type in working schema referenced by the evaluated type reference *reference* no longer exists as a result of a call to PROCESS\_SET\_WORKING\_SCHEMA.

TYPE\_REFERENCE\_IS\_UNSET (*reference*) The type reference *reference* has never been set or has been explicitly unset by TYPE\_REFERENCE\_UNSET.

UNLOCKING\_IN\_TRANSACTION\_IS\_FORBIDDEN To perform the operation would unset or reset a lock within a transaction.

UPPER\_BOUND\_WOULD\_BE\_VIOLATED (*object*, *link*) The creation of *link* from its origin *object* would violate the upper bound of the link type of *link*.

USAGE\_MODE\_ON\_ATTRIBUTE\_TYPE\_WOULD\_BE\_VIOLATED (*object*, [*link*], *attribute*, *mode\_value*) The usage mode of the type in working schema of the attribute *attribute* of the object *object* or of the link *link* of the object *object* does not include *mode\_value*.

USAGE\_MODE\_ON\_LINK\_TYPE\_WOULD\_BE\_VIOLATED (*object*, *link*, *mode\_value*) The usage mode of the type in working schema of the link *link* of the object *object* does not include *mode\_value*.

USAGE\_MODE\_ON\_OBJECT\_TYPE\_WOULD\_BE\_VIOLATED (*initial\_type*, *type*) The usage mode of *initial\_type*, *type*, or any type in working schema which is a descendant of *initial\_type* and an ancestor of *type* does not include CREATE\_MODE.

USER\_CRITERION\_IS\_NOT\_SELECTED (*criterion*) The event to be removed in *criterion* is not in the list of user-dependent criteria.

USER\_GROUP\_IS\_IN\_USE (*user*, *group*) A process executing on behalf of *user* currently has the group *group* as its adopted group.

USER\_GROUP\_LACKS\_ALL\_USERS\_AS\_SUPERGROUP (*group*) The predefined user group ALL\_USERS is not a supergroup of *group*.

**USER\_GROUP\_WOULD\_NOT\_HAVE\_ALL\_USERS\_AS\_SUPERGROUP** (*group*) Deleting the "user\_subgroup\_of" link from *group* would mean that the predefined group ALL\_USERS would no longer be a supergroup of *group*; and either *group* has users as members, or there is a group G which has members and which has *group* as one of its supergroups and G would no longer have ALL\_USERS as a supergroup.

**USER\_IS\_ALREADY\_CLEARED\_TO\_CLASS** (*group*, *class*) The user group *group* is already cleared to a class which dominates or is dominated by *class*.

**USER\_IS\_ALREADY\_MEMBER** (*user*, *group*) The user *user* is already a member of *group*.

**USER\_IS\_IN\_USE** (*user*) The user *user* is the destination of a "user\_identity" link from a process.

**USER\_IS\_NOT\_CLEARED** (*process*, *mandatory\_label*) The user group on whose behalf *process* is executing does not have an overall clearance to the level given by the label *mandatory\_label*. More precisely:

LABEL\_DOMINATES (user confidentiality clearance of *process*, *mandatory\_label*)

**USER\_IS\_NOT\_CLEARED\_TO\_CLASS** (*group*, *class*) The user group *group* is not cleared to mandatory class *class* nor to a mandatory class dominating *class*.

**USER\_IS\_NOT\_MEMBER** (*user*, *user\_group*) *user* is not a member of the user group *user\_group*.

**USER\_IS\_UNKNOWN** (*object*) The user designated by *object* in a user-dependent criterion is not a known user, i.e. it is not a known security group which is of the child type "user" nor of one of its descendent types.

**VALUE\_TYPE\_IS\_INVALID** (*value*, *object*, [*link*], *attribute*) The type of *value* is not of the value type of the associated type of the attribute *attribute* of the object *object* or of the link *link* of the object *object*.

**VERSION\_GRAPH\_IS\_INVALID** (*version*, *predecessor*) Adding *predecessor* as a predecessor of *version* would render its version graph invalid; i.e. not a directed acyclic graph.

**VERSION\_IS\_REQUIRED** (*version*, *scope*) One of the objects specified by *version* and *scope* cannot be removed as it has no successors and there are no incoming composition or existence links other than the predecessor link from one of the objects to at least one of its predecessors.

**VOLUME\_CANNOT\_BE\_MOUNTED\_ON\_DEVICE** (*volume*, *device*) The device *device* is inappropriate for the mounting of the volume *volume*.

**VOLUME\_EXISTS** (*volume\_identifier*) The specified volume number *volume\_identifier* corresponds to an existing volume.

**VOLUME\_HAS\_OBJECT\_OUTSIDE\_RANGE** (*volume*, *high\_label*, *low\_label*) The volume *volume* has one or more objects whose mandatory label lies outside the range *low\_label* to *high\_label*.

**VOLUME\_HAS\_OBJECTS\_IN\_USE** (*volume\_identifier*) The volume *volume\_identifier* cannot be unmounted. This could occur if:

- some resources (objects or links) residing on that volume are currently locked;

- some "sds", "type\_in\_sds", or "type" objects residing on the volume are currently included in some working schemas;
- a message queue object residing on the volume is associated with a reserved message queue;
- some static contexts residing on the volume are being currently executed or are currently used as an interpreter;
- some process objects residing on the volume are associated with processes currently running or some activity objects residing on the volume are associated with non-terminated activities.

VOLUME\_HAS\_OTHER\_LINKS (*volume*) There are links starting from the volume *volume* which are not the "object\_on\_volume" link to itself, the "mounted\_on" link to the device on which the volume is mounted, or the reverse of the link from the volume directory.

VOLUME\_HAS\_OTHER\_OBJECTS (*volume*) The volume *volume* has other objects residing on it apart from the volume object representing the volume itself.

VOLUME\_IDENTIFIER\_IS\_INVALID (*volume\_identifier*) *volume\_identifier* is not a valid volume number.

VOLUME\_IS\_ADMINISTRATION\_VOLUME (*volume*) The volume *volume* is an administration volume.

VOLUME\_IS\_ALREADY\_COPY\_VOLUME\_OF\_REPLICA\_SET (*replica\_set*, *volume*) The volume *volume* is already a copy volume of the replica set *replica\_set*.

VOLUME\_IS\_ALREADY\_MOUNTED (*volume*) The volume *volume* is already mounted.

VOLUME\_IS\_FULL (*volume*) There is insufficient space for one of the objects or for one of the links to be created on the volume *volume*.

VOLUME\_IS\_INACCESSIBLE (*volume*) The volume *volume* is not accessible (see 18.1.5).

VOLUME\_IS\_MASTER\_VOLUME\_OF\_REPLICA\_SET (*replica\_set*, *volume*) The volume *volume* is already the master volume of the replica set *replica\_set*.

VOLUME\_IS\_NOT\_COPY\_VOLUME\_OF\_REPLICA\_SET (*replica\_set*, *volume*) The volume *volume* is not a copy volume of the replica set *replica\_set*.

VOLUME\_IS\_NOT\_MASTER\_OR\_COPY\_VOLUME\_OF\_REPLICA\_SET (*replica\_set*, *volume*) The volume *volume* is neither the master nor a copy volume of the replica set *replica\_set*.

VOLUME\_IS\_READ\_ONLY (*volume*) The volume *volume* is mounted as a read-only volume.

VOLUME\_IS\_UNKNOWN (*volume*) The "volume" object *volume* is not linked to the volume directory.

WORKSTATION\_EXISTS (*identifier*) A workstation already exists in the specified administration volume with the execution site identifier *identifier*.

WORKSTATION\_HAS\_NO\_CHOICE\_OF\_VOLUME\_FOR\_REPLICA\_SET (*station*, *replica\_set*) The workstation *station* has no chosen volume for accessing replica set *replica\_set*.

WORKSTATION\_IDENTIFIER\_IS\_INVALID (*identifier*) The natural *identifier* is not a valid execution site identifier.

WORKSTATION\_IS\_BUSY (*station*) The workstation *station* is busy.

WORKSTATION\_IS\_CONNECTED (*station*) The connection status of the workstation *station* is not LOCAL or AVAILABLE.

WORKSTATION\_IS\_NOT\_CONNECTED (*station*) The status of the workstation *station* is LOCAL or AVAILABLE.

WORKSTATION\_IS\_UNKNOWN (*station*) The workstation *station* is unknown to the PCTE installation (i.e. it is not linked to the execution site directory with a "known\_execution\_site" link).

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC 13719-1:1998

## **Annex D**

(normative)

### **Auditable Events**

#### **D.1 Selectable events**

##### **D.1.1 Selectable event type = WRITE**

ACCOUNTING\_LOG\_COPY\_AND\_RESET  
ACCOUNTING\_OFF  
ACCOUNTING\_ON  
ACCOUNTING\_RECORD\_WRITE  
AUDIT\_FILE\_COPY\_AND\_RESET  
CONSUMER\_GROUP\_DELETE  
CONSUMER\_GROUP\_INITIALIZE  
CONTENTS\_COPY\_FROM\_FOREIGN\_SYSTEM  
CONTENTS\_OPEN  
CONTENTS\_SET\_POSITION  
CONTENTS\_SET\_PROPERTIES  
DEVICE\_CREATE  
DEVICE\_REMOVE  
LINK\_CREATE  
LINK\_DELETE  
LINK\_DELETE\_ATTRIBUTE  
LINK\_REPLACE  
LINK\_RESET\_ATTRIBUTE  
LINK\_SET\_ATTRIBUTE  
LINK\_SET\_SEVERAL\_ATTRIBUTES  
MESSAGE\_SEND\_NO\_WAIT  
MESSAGE\_SEND\_WAIT  
NOTIFY\_DELETE  
NOTIFY\_SWITCH\_EVENTS  
OBJECT\_CONVERT  
OBJECT\_CREATE  
OBJECT\_DELETE  
OBJECT\_DELETE\_ATTRIBUTE  
OBJECT\_RESET\_ATTRIBUTE  
OBJECT\_SET\_ATTRIBUTE  
OBJECT\_SET\_PREFERENCE  
OBJECT\_SET\_SEVERAL\_ATTRIBUTES  
OBJECT\_SET\_TIME\_ATTRIBUTES  
PROCESS\_SET\_ADOPTABLE\_FOR\_CHILD  
PROCESS\_SET\_ALARM  
PROCESS\_SET\_CONSUMER\_IDENTITY  
PROCESS\_SET\_DEFAULT\_ACL\_ENTRY  
PROCESS\_SET\_DEFAULT\_OWNER

PROCESS\_SET\_OPERATION\_TIME\_OUT  
PROCESS\_UNSET\_CONSUMER\_IDENTITY  
QUEUE\_SET\_TOTAL\_SPACE  
RESOURCE\_GROUP\_ADD\_OBJECT  
RESOURCE\_GROUP\_DELETE  
RESOURCE\_GROUP\_INITIALIZE  
RESOURCE\_GROUP\_REMOVE\_OBJECT  
SDS\_ADD\_DESTINATION  
SDS\_APPLY\_ATTRIBUTE\_TYPE  
SDS\_APPLY\_LINK\_TYPE  
SDS\_CREATE\_BOOLEAN\_ATTRIBUTE\_TYPE  
SDS\_CREATE\_ENUMERATION\_ATTRIBUTE\_TYPE  
SDS\_CREATE\_ENUMERAL\_TYPE  
SDS\_CREATE\_FLOAT\_ATTRIBUTE\_TYPE  
SDS\_CREATE\_INTEGER\_ATTRIBUTE\_TYPE  
SDS\_CREATE\_LINK\_TYPE  
SDS\_CREATE\_NATURAL\_ATTRIBUTE\_TYPE  
SDS\_CREATE\_OBJECT\_TYPE  
SDS\_CREATE\_RELATIONSHIP\_TYPE  
SDS\_CREATE\_STRING\_ATTRIBUTE\_TYPE  
SDS\_CREATE\_TIME\_ATTRIBUTE\_TYPE  
SDS\_DELETE\_TYPE  
SDS\_INITIALIZE  
SDS\_REMOVE  
SDS\_REMOVE\_DESTINATION  
SDS\_SET\_ENUMERAL\_TYPE\_IMAGE  
SDS\_SET\_TYPE\_MODES  
SDS\_SET\_TYPE\_NAME  
SDS\_UNAPPLY\_ATTRIBUTE\_TYPE  
SDS\_UNAPPLY\_LINK\_TYPE  
VERSION\_ADD\_PREDECESSOR  
VERSION\_REMOVE  
VERSION\_REMOVE\_PREDECESSOR  
VOLUME\_CREATE  
VOLUME\_DELETE

#### **D.1.2 Selectable event type = READ**

ACCOUNTING\_LOG\_READ  
CONTENTS\_COPY\_TO\_FOREIGN\_SYSTEM  
CONTENTS\_GET\_POSITION  
CONTENTS\_OPEN  
GROUP\_GET\_IDENTIFIER  
LINK\_GET\_ATTRIBUTE  
LINK\_GET\_DESTINATION\_ARCHIVE  
LINK\_GET\_DESTINATION\_VOLUME  
LINK\_GET\_KEY

LINK\_GET\_SEVERAL\_ATTRIBUTES  
MESSAGE\_DELETE  
MESSAGE\_PEEK  
MESSAGE\_RECEIVE\_NO\_WAIT  
MESSAGE\_RECEIVE\_WAIT  
OBJECT\_CHECK\_PERMISSION  
OBJECT\_CHECK\_TYPE  
OBJECT\_GET\_ACL  
OBJECT\_GET\_ATTRIBUTE  
OBJECT\_GET\_PREFERENCE  
OBJECT\_GET\_SEVERAL\_ATTRIBUTES  
OBJECT\_GET\_TYPE  
OBJECT\_LIST\_LINKS  
OBJECT\_LIST\_VOLUMES  
PROCESS\_ADD\_BREAKPOINT  
PROCESS\_CONTINUE  
PROCESS\_GET\_WORKING\_SCHEMA  
PROCESS\_PEEK  
PROCESS\_POKE  
PROCESS\_PROFILING\_OFF  
PROCESS\_PROFILING\_ON  
PROCESS\_REMOVE\_BREAKPOINT  
QUEUE\_EMPTY  
QUEUE\_RESERVE  
QUEUE\_UNRESERVE  
SDS\_GET\_ATTRIBUTE\_TYPE\_PROPERTIES  
SDS\_GET\_ENUMERAL\_TYPE\_IMAGE  
SDS\_GET\_ENUMERAL\_TYPE\_POSITION  
SDS\_GET\_LINK\_TYPE\_PROPERTIES  
SDS\_GET\_NAME  
SDS\_GET\_OBJECT\_TYPE\_PROPERTIES  
SDS\_GET\_TYPE\_KIND  
SDS\_GET\_TYPE\_MODES  
SDS\_GET\_TYPE\_NAME  
SDS\_SCAN\_ATTRIBUTE\_TYPE  
SDS\_SCAN\_ENUMERAL\_TYPE  
SDS\_SCAN\_LINK\_TYPE  
SDS\_SCAN\_OBJECT\_TYPE  
SDS\_SCAN\_TYPES  
VERSION\_IS\_CHANGED  
VERSION\_TEST\_ANCESTRY  
VERSION\_TEST\_DESCENT  
VOLUME\_LIST\_OBJECTS



**D.1.3 Selectable event type = COPY**

ACCOUNTING\_LOG\_COPY\_AND\_RESET  
OBJECT\_COPY  
QUEUE\_RESTORE  
QUEUE\_SAVE  
SDS\_IMPORT\_ATTRIBUTE\_TYPE  
SDS\_IMPORT\_ENUMERAL\_TYPE  
SDS\_IMPORT\_LINK\_TYPE  
SDS\_IMPORT\_OBJECT\_TYPE  
VERSION\_REVERSE  
VERSION\_SNAPSHOT

**D.1.4 Selectable event type = ACCESS\_CONTENTS**

CONTENTS\_READ  
CONTENTS\_TRUNCATE  
CONTENTS\_WRITE

**D.1.5 Selectable event type = EXPLOIT**

PROCESS\_CREATE  
PROCESS\_CREATE\_AND\_START  
PROCESS\_SET\_WORKING\_SCHEMA  
PROCESS\_START  
VOLUME\_MOUNT  
VOLUME\_UNMOUNT

**D.1.6 Selectable event type = CHANGE\_ACCESS\_CONTROL**

OBJECT\_SET\_ACL\_ENTRY

**D.1.7 Selectable event type = CHANGE\_LABEL**

DEVICE\_SET\_CONFIDENTIALITY\_RANGE  
DEVICE\_SET\_INTEGRITY\_RANGE  
EXECUTION\_SITE\_SET\_CONFIDENTIALITY\_RANGE  
EXECUTION\_SITE\_SET\_INTEGRITY\_RANGE  
OBJECT\_SET\_CONFIDENTIALITY\_LABEL  
OBJECT\_SET\_INTEGRITY\_LABEL  
PROCESS\_SET\_CONFIDENTIALITY\_LABEL  
PROCESS\_SET\_FLOATING\_CONFIDENTIALITY\_LEVEL  
PROCESS\_SET\_FLOATING\_INTEGRITY\_LEVEL  
PROCESS\_SET\_INTEGRITY\_LABEL  
VOLUME\_SET\_CONFIDENTIALITY\_RANGE  
VOLUME\_SET\_INTEGRITY\_RANGE

**D.1.8 Selectable event type = VIOLATION\_CONFIDENTIALITY\_WRITE and VIOLATION\_INTEGRITY\_WRITE**

All operations with mandatory security errors plus:

DEVICE\_SET\_CONTROL  
PROCESS\_INTERRUPT\_OPERATION

**D.1.9 Selectable event type = VIOLATION\_CONFIDENTIALITY\_READ and VIOLATION\_INTEGRITY\_READ**

All operations with mandatory security errors plus:

CONTENTS\_GET\_HANDLE\_FROM\_KEY  
CONTENTS\_GET\_KEY\_FROM\_HANDLE  
CONTENTS\_OPEN  
CONTENTS\_SEEK  
DEVICE\_GET\_CONTROL  
LINK\_GET\_REVERSE  
MESSAGE\_RECEIVE\_NO\_WAIT  
MESSAGE\_RECEIVE\_WAIT  
MESSAGE\_SEND\_NO\_WAIT  
MESSAGE\_SEND\_WAIT  
NOTIFY\_CREATE  
OBJECT\_IS\_COMPONENT  
PROCESS\_SET\_REFERENCED\_OBJECT  
PROCESS\_UNSET\_REFERENCED\_OBJECT  
QUEUE\_HANDLER\_DISABLE  
QUEUE\_HANDLER\_ENABLE  
VOLUME\_GET\_STATUS

**D.1.10 Selectable event type = USE\_PREDEFINED\_GROUP**

All uses of type identifiers for access to non-visible types (PCTE\_CONFIGURATION)

All modifications to master objects (PCTE\_REPLICATION)

All operations defined in 10.2

ARCHIVE\_RESTORE  
ARCHIVE\_SAVE  
DEVICE\_CREATE  
DEVICE\_REMOVE  
OBJECT\_SET\_TIME\_ATTRIBUTES  
PROCESS\_SET\_FILE\_SIZE\_LIMIT  
PROCESS\_SET\_PRIORITY  
REPLICA\_SET\_CREATE  
REPLICA\_SET\_REMOVE  
REPLICA\_SET\_ADD\_COPY\_VOLUME  
REPLICA\_SET\_REMOVE\_COPY\_VOLUME  
REPLICATED\_OBJECT\_CREATE

REPLICATED\_OBJECT\_DELETE\_REPLICA  
REPLICATED\_OBJECT\_DUPLICATE  
REPLICATED\_OBJECT\_REMOVE  
TIME\_SET  
VERSION\_ADD\_PREDECESSOR  
VERSION\_REMOVE  
VERSION\_REMOVE\_PREDECESSOR  
VOLUME\_CREATE  
VOLUME\_DELETE  
WORKSTATION\_REDUCE\_CONNECTION  
WORKSTATION\_CONNECT  
WORKSTATION\_CREATE  
WORKSTATION\_DELETE  
WORKSTATION\_DISCONNECT  
WORKSTATION\_SELECT\_REPLICA\_SET\_VOLUME  
WORKSTATION\_UNSELECT\_REPLICA\_SET\_VOLUME

**D.1.11 Selectable event type = SET\_USER\_IDENTITY**

PROCESS\_ADOPT\_USER\_GROUP  
PROCESS\_SET\_USER

**D.1.12 Selectable event type = COVERT\_CHANNEL**

ACTIVITY\_ABORT  
ACTIVITY\_END  
ACTIVITY\_START  
CONTENTS\_GET\_POSITION  
CONTENTS\_READ  
CONTENTS\_SET\_POSITION  
CONTENTS\_TRUNCATE  
CONTENTS\_WRITE  
LINK\_CREATE  
LOCK\_RESET\_INTERNAL\_MODE  
LOCK\_SET\_INTERNAL\_MODE  
LOCK\_SET\_OBJECT  
LOCK\_UNSET\_OBJECT  
PROCESS\_RESUME  
PROCESS\_SET\_REFERENCED\_OBJECT  
PROCESS\_SUSPEND  
PROCESS\_TERMINATE

**D.2 Mandatory Events**

**D.2.1 Mandatory event type = CHANGE\_IDENTIFICATION**

OBJECT\_MOVE

**D.2.2 Mandatory event type = SELECT\_AUDIT\_EVENT**

AUDIT\_ADD\_CRITERION  
AUDIT\_FILE\_COPY\_AND\_RESET  
AUDIT\_REMOVE\_CRITERION  
AUDIT\_SELECTION\_CLEAR  
AUDIT\_SWITCH\_OFF\_SELECTION  
AUDIT\_SWITCH\_ON\_SELECTION

**D.2.3 Mandatory event type = SECURITY\_ADMINISTRATION**

CONFIDENTIALITY\_CLASS\_INITIALIZE  
GROUP\_DELETE  
GROUP\_DISABLE\_FOR\_CONFIDENTIALITY\_DOWNGRADE  
GROUP\_DISABLE\_FOR\_INTEGRITY\_UPGRADE  
GROUP\_ENABLE\_FOR\_CONFIDENTIALITY\_DOWNGRADE  
GROUP\_ENABLE\_FOR\_INTEGRITY\_UPGRADE  
GROUP\_INITIALIZE  
GROUP\_RESTORE  
INTEGRITY\_CLASS\_INITIALIZE  
PROGRAM\_GROUP\_ADD\_MEMBER  
PROGRAM\_GROUP\_ADD\_SUBGROUP  
PROGRAM\_GROUP\_REMOVE\_MEMBER  
PROGRAM\_GROUP\_REMOVE\_SUBGROUP  
USER\_EXTEND\_CONFIDENTIALITY\_CLEARANCE  
USER\_EXTEND\_INTEGRITY\_CLEARANCE  
USER\_GROUP\_ADD\_MEMBER  
USER\_GROUP\_ADD\_SUBGROUP  
USER\_GROUP\_REMOVE\_MEMBER  
USER\_GROUP\_REMOVE\_SUBGROUP  
USER\_REDUCE\_CONFIDENTIALITY\_CLEARANCE  
USER\_REDUCE\_INTEGRITY\_CLEARANCE

**D.3 List of operations not audited**

AUDIT\_FILE\_READ  
AUDIT\_GET\_CRITERIA  
AUDIT\_RECORD\_WRITE  
AUDITING\_STATUS  
CONTENTS\_CLOSE  
NOTIFICATION\_MESSAGE\_GET\_KEY  
PROCESS\_GET\_DEFAULT\_ACL  
PROCESS\_GET\_DEFAULT\_OWNER  
PROCESS\_SET\_TERMINATION\_STATUS  
PROCESS\_WAIT\_FOR\_ANY\_CHILD  
PROCESS\_WAIT\_FOR\_CHILD  
REFERENCE\_COPY  
REFERENCE\_GET\_EVALUATION\_POINT