General DDML document format

The general format of a DDML document looks as follows:

Tag	Description
`<tables-definition>`	Root element of a DDML document. Its children are the main database entities: `<table-definition>`, `<ddl-definition>`, and `<view-definition>`. The `<tables-definition>` tag includes the product attribute, which indicates the product of this specific DDML definition.
`<table-definition>`	Includes the table-level attributes. Its children are the following tags: • `<column-definition>` • `<index-definition>` • `<foreign-key-definition>` If a table is a statistics table, meaning that the type attribute has the value `STATISTICS`, the following tags can be specified as children of the `<table-definition>` tag: • `<summary-hour>` • `<summary-day>` • `<summary-week>` • `<summary-month>`
`<ddl-definition>`	Includes ddl-level attributes.
`<view-definition>`	Includes view-level attributes.

<tables-definition> tag

Root element of the DDML document.

Attribute	Definition
`product`	*Required*. Holds the product short name (product shortcut) consisting if two characters that define the product. In this case, three custom-defined products are available: • `C1` For customer-defined 1 • `C2` For customer-defined 2 • `C3` For customer-defined 3

<table-definition> tag

Holds all of the information included in a single table.

Attribute	Definition
`name`	*Required*. Holds the table name.
`type`	*Required*. Specifies the table type and can hold one of the following values: • **Statistics**. A PMDB table that holds information about certain database activity, summarized into hours or time slices. • **Events**. A PMDB table that keeps a log of database incidents.
`pctfree`	*Optional*. Affects only Oracle and is developed as `PCTFREE` *`pctfree`* . Can have a value between 0 and 99. If the PMDB is a Microsoft SQL Server database, this attribute is transformed to `fill factor`.
`pctused`	*Optional*. Affects only Oracle and is developed as `PCTUSED pctused`. Can have a value between 0 and 99.
`initrans`	*Optional*. Affects only Oracle and is developed as `INITRANS initrans`. Can have a value between 1 and 255.
`maxtrans`	*Optional*. Affects only Oracle and is developed as `MAXTRANS maxtrans`. Can have a value between 1 and 255.

`oracle-storage-clause`	*Optional*. Affects only Oracle and is developed as `STORAGE(oracle-storage-clause)`. For example: `oracle-storage-clause="initial 1M next 1M minextents 1maxextents unlimited pctincrease 100"`
`oracle-additional-clause`	*Optional*. Added to support any other Oracle features that cannot be defined in an Oracle storage clause. For example: `NOLOGGING`
`filterable`	Required for statistics tables (`type=STATISTICS`). If one or more columns are not relevant or needed in the PMDB, set this attribute to `TRUE`. If a statistics table is filterable and some columns are specified as not needed in the load or summary control files, these columns are not loaded or summarized.

The `<table-definition>` tag can have the following children:

- `<column-definition>`
- `<index-definition>`
- `<foreign-key-definition>`

If a table is a statistics table, meaning that the type attribute has the value `STATISTICS`, the following tags can be specified as children of the `<table-definition>` tag:

Additional Child	Definition
`<summary-hour>`	*Optional*. Instructs to define an hour-level summary table for this table. Used for time slice statistics tables. The name of the summary table is specified by the name attribute (*required*). The summary table name should follow the naming conventions specified in Naming conventions.
`<summary-day>`	*Optional*. Instructs to define a day-level summary table for this table. The name of the summary table is specified by the name attribute (*required*). The summary table name should follow the naming conventions specified in Naming conventions.
`<summary-week>`	*Optional*. Instructs to define a week-level summary table for this table. The name of the summary table is specified by the name attribute (*required*). The summary table name should follow the naming conventions specified in Naming conventions.
`<summary-month>`	*Optional*. Instructs to define a month-level summary table for this table. The name of the summary table is specified by the name attribute (*required*). The summary table name should follow the naming conventions specified in Naming conventions.

<column-definition> tag

Holds all the parameters of a column.

Attribute	Definition
`name`	*Required*. Holds the column name. The name should follow the naming conventions specified in Naming conventions.

`data-type`	*Required*. Holds the data type of a column:
	
	• `BOOLEAN` Does not require `data-length` or `data-scale` attributes. These attributes should not be specified. Implementation on all DBMSs:
	  ○ `CHAR(1)` where `T` is the boolean `TRUE` value and `F` is `FALSE`.
	• `CHAR` Requires `data-length` attribute. Implementation on all DBMSs:
	  ○ `CHAR(`*`data-length`*`)`
	• `CLOB` Requires `data-length` attribute. Implementation:
	  ○ `TEXT` for Microsoft SQL Server
	  ○ `CLOB` for Oracle
	• `DECIMAL` Requires `data-length` (used for precision) and `data-scale` attributes. Implementation:
	  ○ `NUMBER(`*`data-length, data-scale`*`)` for Microsoft SQL Server and IBM UDB Oracle
	• `FLOAT` Requires `data-length` attribute. Implementation:
	  ○ `FLOAT(`*`data-length`*`)` for Microsoft SQL Server
	  ○ `NUMBER` without any parameters for Oracle
	• `INTEGER` Requires `data-length` attribute. Implementation:
	  ○ `data-length` less than three: `TINYINT`; `data-length` two-to-four: `SMALLINT`; `data-length` more than four: `BIGINT` for Microsoft SQL Server
	  ○ `NUMBER(`*`data-length`*`)` for Oracle
	• `TIMESTAMP` Does not require `data-length` or `data-scale` attributes. These attributes should not be specified. Implementation:
	  ○ `DATETIME` for Microsoft SQL Server
	  ○ `TIME` for Oracle
	• `UNIQUE INTEGER` Does not require `data-length` or `data-scale` attributes. These attributes should not be specified. This is a data type for a unique integer whose values are generated automatically. Implementation:
	  ○ `IDENTITY` for Microsoft SQL Server
	  ○ `NUMBER(20,0)` for Oracle. In addition, a `SEQUENCE` and a TRIGGER BEFORE INSERT are created, which select the `NEXTVAL` of the `SEQUENCE` from `DUAL`.
	• `VARCHAR` Requires the `data-length` attribute. Implementation:
	  ○ `VARCHAR(`*`data-length`*`)` for Microsoft SQL Server
	  ○ `VARCHAR2(`*`data-length`*`)` for Oracle
	• `VARBINARY` Requires the `data-length` attribute. Implementation:
	  ○ `VARBINARY(data-length)` for Microsoft SQL Server
	  ○ `RAW(data-length)` for Oracle

`data-length`	Required only for the data types listed below. Specifies the column data length: • `CHAR` • `CLOB` • `DECIMAL` • `FLOAT` • `INTEGER` • `VARCHAR` • `VARBINARY`
`data-scale`	Required only for the data type `DECIMAL`. Specifies the column data scale.
`null`	*Required*. Has the value `TRUE` if the column is nullable and `FALSE` if it is not.

`default`	*Optional*. Specifies the column default. Can have the following values:
	
	• `NULL` if the column is nullable (null=TRUE).
	• `N/A` if no default exists. If the attribute is omitted, this is the default.
	
	A constant value that is equal to the column type:
	
	• `BOOLEAN` `TRUE` or `FALSE`
	• `CHAR` A textual constant, such as `ABC` `CLOB`; a textual constant, such as `ABC` `DECIMAL`; or a decimal point constant, such as `10.3`
	• `FLOAT` A floating point constant, such as `12E7`
	• `INTEGER` An integer constant, such as `27`
	• `TIMESTAMP` One of the following:
	  ○ A timestamp constant of the format `yyyy-mm-dd hh:mm:ss.fffffffff`, which is jafa.sql.Timestamp's format, such as `2020-02-20 23:07:35:175000000`. Each DBMS displays a slightly different default. Implementation:
	  ▪ `2020-02-20 23:07:35:175` for Microsoft SQL Server
	  ▪ `TO_DATE('2020-02-20 23:07:35')` for Oracle
	  ○ The literal string `CURRENT TIMESTAMP`. Implementation:
	  ▪ `GETDATE()` for Microsoft SQL Server
	  ▪ `SYSDATE` for Oracle
	• `UNIQUE INTEGER` Default value not required and not allowed.
	• `VARCHAR` A textual constant, such as `ABC`
	• `VARBINARY` A hex string where every two hexadecimal digits represent one byte, such as `A07C889F`. Each DBMS displays a slightly different default. Implementation:
	  ○ `0xA07C889F` for Microsoft SQL Server
	  ○ `HEXTORAW('A07C889F')` for Oracle

`type`	Required for the column role in columns of statistics tables. Can have the following values: • `IDENTIFIER` A column identifying the sampled entity. The concatenation of all identifiers should uniquely identify the entity. • `DATE` A column identifying the sampled period. Its type should be `TIMESTAMP` and be equal to the beginning of the sampled period. • `SUM` A statistics column whose transfer to a higher summary level (such as hourly to daily) should be applied by the `SUM` function. • `MIN` A statistics column whose transfer to a higher summary level (such as hourly to daily) should be applied by the `MIN` function. • `MAX` A statistics column whose transfer to a higher summary level (such as hourly to daily) should be applied by the `MAX` function. • `AVG` A statistics column whose transfer to a higher summary level (such as hourly to daily) should be applied by the `AVG` function.

Columns required for statistics tables

If the table is a statistics table, you must include the following columns:

- `<column-definition name="<table shortcut>_TIMESTAMP" data-type="TIMESTAMP" null="FALSE" type="DATE"/>`
- `<column-definition name="<table shortcut>_MINUTES_COUNT_SUM" data-type="INTEGER" data-length="9" null="FALSE" type="SUM" default="1"/>`
- `<column-definition name="<table shortcut>_PWHG_ID" data-type="INTEGER" data-length="4" null="FALSE" type="IDENTIFIER"/>`
- `<column-definition name="<table shortcut>_PWII_INSTANCE_ID" data-type="INTEGER" data-length="9" null="FALSE" type="IDENTIFIER"/>`
- `<column-definition name="<table shortcut>_RECIEVED_TIMESTAMP" data-type="TIMESTAMP" null="FALSE" default="CURRENT TIMESTAMP" type="DATE"/>`

> You must replace `<table shortcut>` with the four characters that represent the relevant table (see Table Shortcut).

<index-definition> tag

Holds all the parameters for an index definition.

Attribute	Definition
`name`	*Required*. Holds the index name. The name should follow the naming conventions specified in Naming conventions.
`unique`	*Required*. Has the value `TRUE` if the index is unique and `FALSE` if it is not.
`primary`	*Required*. Has the value `TRUE` if this is a primary index and `FALSE` if it is not. Implementation: • `ALTER TABLE ADD CONSTRAINT` for Microsoft SQL Server. Adding a primary constraint in Microsoft SQL Server always results in creating a unique index to enforce the constraint. • An index is created using the `CREATE INDEX` statement for Oracle. Then an `ALTER TABLE ADD CONSTRAINT` is performed to add the primary constraint. The `ALTER TABLE ADD CONSTRAINT` is suffixed with the `USING INDEX` clause to instruct Oracle to use the already created index to enforce the constraint and not to create a new one.
`clustered`	*Optional*. Has the value `TRUE` if this is a clustered index and `FALSE` if it is not. The default is `FALSE` . A clustered index is an index that physically orders and organizes the table. Implementation: • `CLUSTERED` clause for a clustered index, `NON CLUSTERED` clause for a non-clustered index in Microsoft SQL Server. • In Oracle, this feature is not used frequently; it requires a complete entity to handle clustering.
`mssql-additional-clause`	*Optional*. Only relevant for Microsoft SQL Server. Allows specifying every parameter defined in the "with" section.

<foreign-key-definition> tag

Holds all the parameters for a foreign-key definition.

Attribute	Definition
`name`	*Required*. Holds the foreign key constraint name. The name should follow the naming conventions specified in Naming conventions.
`ref-table`	Required for the table name referenced by the foreign key.
`ref-columns`	*Required*. Includes pairs of referencing and referenced columns separated by blanks.

`on-delete`	*Required*. Determines what action is taken if one or more rows in the referencing table point to a row in the referenced table that has been deleted. Can have the following values:
	• `CASCADE` All rows pointing to the deleted row are also deleted. Implementation on all DBMSs: `ON DELETE CASCADE`
	• `NO ACTION` The deletion fails. Implementation:
	  ○ `ON DELETE NO ACTION` for Microsoft SQL Server
	  ○ No `ON DELETE` clause is specified for Oracle. This is the default.

`on-update`	*Required*. Determines what action is taken if one or more rows in the referencing table point to a row in the referenced table that has been updated. Can have the following values:
	• `CASCADE` All rows pointing to the deleted row are also deleted. Implementation on all DBMSs: `ON UPDATE CASCADE`
	• `NO ACTION` The update fails. Implementation:
	  ○ `ON UPDATE NO ACTION` for Microsoft SQL Server
	  ○ No `ON UPDATE` clause is specified for Oracle. This is the default.

<ddl-definition> tag

Holds DDL (Data Definition Language) and DML (Data Manipulation Language) statements that can be performed during installation or uninstallation, such as stored procedures.

Attribute	Definition
`statement`	*Required*. Defines the DDL statement.
`dbms`	*Optional*. the RDBMS type on which the DDL generates. Must be one of the following: • `oracle` • `mssql` • `""` If left empty (`""`), the DDL is created on all relational database management systems (RDBMS).
`version`	*Optional*. The RDBMS version on which the DDL generates. Should have the format `8.1. ...` If left empty, the DDL is created on all RDBMS versions.
`event`	*Optional*. The RDBMS version on which the DDL generates. Must be one of the following: • `INSTALL` *Default*. During the installation process only. • `UNINSTALL` During the uninstallation process only. • `PREINSTALL` Before the installation process. • `PREUNINSTALL` Before the uninstallation process.

<view-definition> tag

Holds all the parameters required for a view definition.

Attribute	Definition
`name`	*Required*. Holds the view name. The name should follow the naming conventions specified in Naming conventions.
`type`	*Required*. Specifies the view type and can hold one of the following values: • `STATISTICS` A PMDB view that holds information about certain database activity, summarized into hours or time slices. • `INTERNAL` A PMDB view that keeps any other, non-statistical information.
`view-columns`	*Required*. Its value is in the view columns separated by blanks. The number of columns should be equal to the number selected in the `as-query` attribute.
`as-query`	*Required*. Specifies the select table that defines the view.
`check-option`	*Optional*. Default value is `FALSE`. If the value is `TRUE`, only modifications that are visible through the view are allowed, meaning that `INSERT` and `UPDATE` statements are valid only if the affected rows can be retrieved by the view afterward.