

Database Design: A Quality Control Checklist

Relationships

- ❑ There are no many-to-many relationships.
- ❑ If there are any one-to-one relationships, there is a strong justification for them.
- ❑ The entity-relationship diagram shows the cardinality of all the relationships linking any two tables.
- ❑ Every table pointing to another table with a “one” cardinality includes in it the primary key of that table as a foreign key.

Tables

- ❑ Tables have full, descriptive names reflecting the entities/events they are supposed to represent; the names appear in the singular.
- ❑ The master files appear on the outside and the transaction files (whose job is to link master files) on the inside in order to avoid crossing lines.
- ❑ All transaction files have date fields in them. Name each date field uniquely.
- ❑ There are no single-column look-up tables; each look-up table has a dedicated primary key. [The reason is ...](#)

Fields

- ❑ Fields have full, descriptive names reflecting the attributes they are supposed to represent, particularly for unique, application-specific fields. Be prepared to use the same field names in the implementation process.
- ❑ Derived (calculated) fields are not stored. Instead, the historic fields based on which the derived fields can be recalculated in the future are stored in the appropriate transaction files.
- ❑ Foreign keys and the primary keys they represent have identical names.
- ❑ A uniform set of conventions in naming fields is followed consistently. Example: avoid using the hyphen in some fields (*material-name*) while the underscore in others (*material_name*).
- ❑ If any field names have to be abbreviated, those abbreviations are meaningful and easy to guess.
- ❑ Primary and foreign key fields are explicitly designated as such.