

# Tables

Most often, additional information (such as materials, countries, currencies, etc.) will not differ within one case.

Therefore, it is reasonable to split the event log into (at least) two tables:

An **activity table** and a **case table**.

## Activity Table

The activity table usually has the **same structure as a minimal event log** (potentially including a sorting column).

Also, **activity-specific information** can be added to the activity table, such as the user which performed an activity in the IT system.

[blocked URL](#)

*Activity table consisting of minimal event log, sorting and one activity-specific information column.*

## Case Table

The case table may contain **case-specific information** and will be linked to the activity table over a **foreign key relation**.

The Case ID must be a primary key of the case table (i.e. each Case ID appears only once in the case table). Hence, the information is only stored once per case.

This prevents unnecessary data explosion.

[blocked URL](#)

*Case table with three case-specific columns.*

## Metadata

It can make sense to extend the data model with further metadata on your processes' cases. You can simply add those additional tables and join them to the cases via a foreign key relation.