

## Referential Integrity

The concept of referential integrity has to do with the relationships between records in associated tables when a foreign key relationship exists. If referential integrity exists, then foreign key relationships are strictly enforced to determine the association between records in parent/child tables. For each record in the parent table, it is possible to have: a) zero, b) exactly one, c) zero or one, or d) zero, one, or more than one child record. For each child record, there is always a parent record. Database administrators have the ability to strictly enforce referential integrity to ensure that these relationships are always followed.

When referential integrity exists, insertions, updates, and deletions in the database are each checked automatically such that records cannot be inserted, updated, or deleted where such actions violate referential integrity requirements. Systems populating data into the database must therefore be designed with this requirement in mind and users manually populating data must be aware as well<sup>1</sup>.

Unfortunately, referential integrity is not always strictly enforced. When it is not enforced, it is possible that records in the child table will not have associated records in the parent table. In this case, the child record is referred to as an orphan. For insurance and reinsurance data, it is common that referential integrity is NOT enforced and orphan records are the result. The reasons for the lack of referential integrity are rooted in the insurance process itself and the legacy systems supporting it. Insurance and reinsurance are constantly evolving business processes. As companies learn that new data elements can enhance their ability to compete, they add them into their processing systems and ultimately into their databases. Most of the time, it is not possible to back fill all of the historical data with the newly discovered fields. In this case, orphan records are frequently the result. Manual journal records (general ledger) and corrective records in legacy mainframe processing systems historically lacked detail for some fields. Again, this leads to orphan records. When you are working with insurance and reinsurance data, you should always start with the assumption that orphans exist unless either: a) a database administrator confirms that referential integrity is strictly enforced or b) you have tested the data to confirm that referential integrity exists. In practice, I always test, even if a database administrator tells me that referential integrity is strictly enforced.

---

<sup>1</sup> While it is almost never desirable to have users manually populating data, it does sometimes happen.