





## #2 - 06/07/2023 02:37 PM - Eric Faulhaber

Constantin Asofiei wrote:

If two tables have the same DUMP-NAME in the .df, the FWD can not start, as it enforces an unique index in the meta table.

Editing the .df by hand can lead to mistakes where the same name is used. FWD should not abend this way.

That makes sense. When we implemented the infrastructure for static metadata, no consideration was made for hand-edit mistakes in the DF file. It was assumed the input was correct and consistent.

These types of constraint violations should be caught and reported during conversion, because in some cases, there will be no good way to recover from invalid metadata at runtime. In this case, DUMP-NAME is not used directly by the FWD runtime, but it might be used by an application, for data export purposes, for example.

The obvious short term workaround is to correct the duplicate DUMP-NAME value which caused this abnormal end. For a more robust, longer term solution, I think we need a general-purpose conversion-time fix.

Static metadata content is collected by the metaschema.xml rule set. Perhaps we do an early walk which collects the unique index information for each metadata table we support. We might be able to use or adapt the UniqueTracker class to enforce unique constraints during the collection of static metadata, so we catch any errors like this one at conversion time, rather than at runtime.

## #3 - 06/08/2023 06:07 AM - Greg Shah

If the 4GL doesn't care that the DUMP-NAME is unique, why do we care? It seems like there is no compatibility requirement here. We've introduced a problem that doesn't exist in OE.

In other words, if the value is incorrect that should not stop the FWD server from running. I don't see the conversion as a good solution. Hand edits to DMOs later might introduce the same problem. The runtime should not be sensitive to this. And I wouldn't want to kill the conversion over this either.

## #4 - 06/08/2023 06:09 AM - Constantin Asofiei

There is an unique index on the \_field.\_dump-name, in the meta schema:

```
ADD INDEX "_Dump-Name" ON "_File"  
  AREA "Schema Area"  
  UNIQUE  
  INDEX-FIELD "_Db-rcid" ASCENDING  
  INDEX-FIELD "_Dump-name" ASCENDING ABBREVIATED  
  INDEX-FIELD "_Owner" ASCENDING
```

The problem here is with manual edit of the .df file, where things can go wrong.

**#5 - 06/08/2023 06:32 AM - Greg Shah**

Why would the dump-name itself be the cause of a collision?

What happens in OE when this same .df is processed?

**#6 - 06/08/2023 06:36 AM - Constantin Asofiei**

Greg Shah wrote:

Why would the dump-name itself be the cause of a collision?

Well, the index is unique, and I assume \_Db-recid and \_Owner is the same for all records.

What happens in OE when this same .df is processed?

I just tested and OE makes the dump-name unique, I used year1 and got year1-----1.

**#7 - 06/09/2023 10:52 AM - Eric Faulhaber**

We could do something similar for the dump name during conversion, if there is a unique constraint violation, so conversion doesn't stop.

However, that doesn't address the general purpose problem of metadata unique constraint violations due to hand-editing mistakes.