

# Re: Snapshot Delivery – Violation of PRIMARY KEY constraint

---

*Source:*

<http://www.tech-archive.net/Archive/SOL-Server/microsoft.public.sqlserver.replication/2007-05/msg00068.html>

---

- *From:* "Raymond Mak [MSFT]" <[rmak@xxxxxxxxxxxxxxxxxxxxxxxx](mailto:rmak@xxxxxxxxxxxxxxxxxxxxxxxx)>
  - *Date:* Thu, 10 May 2007 10:04:57 -0700
- 

Hi Richard,

The simplest explanation is that the filter you use for deleting data at the subscriber doesn't account for all the rows that will be replaced by the bcp data delivered with the snapshot. This is likely to be the case if the filter that you define has the potential to be evaluated differently at the publisher and at the subscriber. One example will be a filter that uses the getdate() function as the current date when the snapshot was generated is most likely different from the current date when the snapshot was delivered. Another example will be a filter that is based on data from another table.

Delivering a snapshot is fundamentally a very disruptive process at the subscriber database although I have seen many attempts to keep the subscriber data available while a snapshot is being delivered. At best, such attempts will end up causing significant performance degradation in the snapshot delivery process as it would be nearly impossible to fulfill the minimal-logging requirements when the snapshot data is being bulk-loaded. At worst, folks run into difficult to resolve issues such as the one that you encountered.

Perhaps a better strategy for making subscriber data available while snapshot is being delivered is to implement some kind of "double-buffering" mechanism. That is, snapshot data can be first delivered to a staging area while keeping the existing subscriber data available which will then be brought online once the snapshot is fully delivered. One way to do this is to simply setup a different subscriber database on the same server and switch your application to point to the new subscriber database once it is fully initialized. Another way is for me to implement the same at the table level using table rename and/or partition switching. However, given the current lack of resource on my end, you may want to log a feedback item @ <http://connect.microsoft.com> if you ever want to see this implemented. Note that either way you pay for the higher data availability with extra disk space usage (up to 2x) but I guess this is typically how such things go.

Hope that helps,

–Raymond

Re: Snapshot Delivery – Violation of PRIMARY KEY constraint

"Richard T" <RichardT@xxxxxxxxxxxxxxxxxxxxxxxxxxxx> wrote in message  
news:4FABE3E8-A8B7-43C6-A669-0B5AA6EF5105@xxxxxxxxxxxxxxxxxxxx

I'm having a problem with transactional replication. The problem arises when applying a snapshot to a subscriber.

The 'Action if name is in use' property of the articles is set to 'Delete data. If article has a row filter, delete only data that matches the filter'

I filter the articles by use of a function which returns 1 or 0 depending on whether the record is to be replicated.

When the bcp process runs it errors with Violation of PRIMARY KEY constraint 'PK\_Article\_ArticleMaster'. Cannot insert duplicate key in object 'dbo.ArticleMaster'.

If I change the 'Action if name is in use' property of the articles to 'Truncate all data in the existing object' the snapshot is applied successfully.

My problem here is that the database is used as the source of a website and needs to be available 24/7, so if I truncate the tables when applying the snapshot visitors of the website will be presented with an empty screen.

I has assumed that the 'Delete data. If article has a row filter, delete only data that matches the filter' setting would remove those records which are about to be bcp'd into the table, am I wrong?

Does anyone have an explanation of what's happening?

Thanks, Richard