

Re: transaction disabling

Source:

<http://www.tech-archive.net/Archive/SQL-Server/microsoft.public.sqlserver.server/2004-03/3114.html>

From: Andrew J. Kelly (*sqlmvpnoooospam_at_shadhawk.com*)

Date: 03/22/04

Date: Sun, 21 Mar 2004 21:01:13 -0600

Glenn,

Still not 100% sure I understand what your doing but it sounds bad. It sounds like you are changing the rows in the lookup tables including their PK with new values and then need to update the FK's in the larger tables to match. You should never change a PK. If you need to change the description that's OK there is no need to change the PK and that should not necessitate the need to change the parent table. Hopefully I am just not understanding why your doing this. In either case you can usually do the small batches by using a WHILE loop and ROWCOUNT.

```
SET ROWCOUNT 10000
```

```
WHILE 1 = 1  
BEGIN
```

```
    UPDATE YourTable SET FK = x WHERE FK = y
```

```
    IF @@ROWCOUNT = 0  
        BREAK  
END
```

```
SET ROWCOUNT 0
```

```
--  
Andrew J. Kelly  SQL MVP  
"Glenn" <g.e.r@comcast.net> wrote in message  
news:OmqDHi5DEHA.2088@TK2MSFTNGP10.phx.gbl...  
> Andrew:  
>  
> Thanks. All this is internal to the database. The lookup tables are just  
> small tables in the database that a user can change. The keys of these  
> tables I want to add directly to a large table in the database. Every row  
> in  
> the large table will have these fields updated to the joined foreign key  
> fields.  
>  
> I cannot delete the records in the large tables, I just need to update a  
> couple of fields to the new/changed lookup key values. Thus I am running  
> an update sproc joining the large table to the lookup tables.
```

microsoft.public.sqlserver.server: Re: transaction disabling

>
>
> Just FYI, I am doing this as I have a dynamic query building form that users
> can build queries off of many base queries hitting many different large
> tables that join the same lookup tables. The filter keys are higher in the
> hierarchy than the large base tables and for some reason the query optimizer
> is not able propagate down to the large tables and the queries were running
> poorly. By filtering the large table directly by using these stored keys the
> queries go down from 40 minutes to 5 seconds.
>
> How can I limit the number of rows in a transaction in a query? Just run
> the update query in a loop filtering on key ranges?
>
> Best,
> Glenn
>
>
> "Andrew J. Kelly" <sqlmvpnooospan@shadhawk.com> wrote in message
> news:uhx\$RS5DEHA.580@TK2MSFTNGP11.phx.gbl...
> > You can not disable the logging but you can under certain circumstances do
> a
> > minimally logged load. Not really sure what your up against without DDL
> or
> > samples but here are some suggestions. If this is one or a few tables
> that
> > are strictly lookup tables I would create a text file with the values
> you
> > want the table to be first. Then truncate the existing table and BCP
> the
> > file into the table. The BCp should be able to happen in a few minutes
> at
> > most if done correctly. Check out the section in BOL under "minimally
> > logged bulk copy" for the details. If this isn't feasible I would try
> doing
> > the updates in smaller batches of say 10000 rows or so at a time. This
> will
> > keep your log file from filling and should dramatically increase the
> speed
> > of the overall operation. Last I suggest you manually increase the log
> file
> > to larger than you think it needs for the operation before you start.
> Make
> > sure the log file is on a separate physical drive (preferably a RAID 1)
> from
> > the data drive. The growing can be very expensive and time consuming and
> > will only slow down your operation.
>
>
> > --
> > Andrew J. Kelly SQL MVP
>
>
>
> > "Microsoft Public" <g.e.r@comcast.net> wrote in message
> > news:uDDYZ64DEHA.624@TK2MSFTNGP10.phx.gbl...
> > > Is it possible to suspend/disable transaction logging on a query? I
> have
> > a
> > > static database that I need to update a table of @ 5 million rows

microsoft.public.sqlserver.server: Re: transaction disabling

```
> through
> > a
> > > single query every couple of days that updates some static lookup
fields
> > > based on other small tables which may have changed.
> > >
> > > The recovery mode is simple, but the update runs about 3 hours which
> > > seems
> > > ridiculously long even though the server is not that fast (450MH) and
> > > memory
> > > only 750MB, though there is plenty of disk space.
> > >
> > > Is there a better way of doing this, a bulk update mode? The log file
> > > grows
> > > to about 2 GB. I have it set up to grow by 250MB chunks.
> > >
> > > Thanks.
> > >
> > > Best,
> > > Glenn
> > >
> > >
> >
>
>
```