
Sterling Control Center Version 5.2 Database Partitioning

The Need for Database Partitioning

Prior to Sterling Control Center 4.2, only one set of database tables was used. When it was time to perform database maintenance (e.g. archive old data off of the Database, rebuild indexes, etc), Sterling Control Center would have to be shut down. When Sterling Control Center was restarted, it would be in “catch up” processing until all missed records were collected from the servers. Although this process worked well for most monitored servers, there were two major exceptions.

Not all monitored servers can recover missed events. For Connect:Enterprise z/OS, Connect:Direct Select, z/OS FTP and WS_FTP servers, events are not collected via a Sterling Control Center request but via SNMP traps. Any events sent from those servers during the time Sterling Control Center was down were lost.

In some cases “catch up” processing caused excessive overhead on Sterling Control Center and the monitored servers. If the monitored servers were very busy during “Catch up” processing, the overhead of gathering missed events could impact processing performance on the monitored system.

Starting with Sterling Control Center 4.2, the concept of a Production database and a Staging database was introduced. Current data is stored in the Production database tables. Older data is moved from the Production database tables to the Staging database tables in the background, on a periodic (one or more times per day) basis. The Production database tables are always available to accept events either from collected events or received SNMP events. When old data needs to be archived from the Staging database tables, the Staging database tables can become unavailable to Sterling Control Center with no impact to Production database processing. It was recommended that the Production and the Staging database tables be on different machines to enable the best possible throughput and performance. Although this configuration works well for most customers, there were two major exceptions.

Some customers were required to purchase two database licenses.

Some customers experience a problem where the movement of data from the Production to the Staging database tables took too long because of the sheer volume of data or because of limited computer processing capacity.

Sterling Control Center 5.1 made several improvements to database processing including reducing the overhead when moving data from the Production database tables to the Staging database tables. In order to take advantage of this improved processing, the Production database tables and the Staging database tables must be defined in the same database installation and the userid for the Production database must be given authority to the Staging database tables. The new recommendation for database configuration in Sterling Control Center 5.1 is to define both the Production database tables and the Staging database tables on a single database installation.

Use of partitioned database tables is an additional database configuration option for Sterling Control Center version 5.2. Rather than defining a set of Production database tables and a set of Staging database tables, use of daily Partitioned database tables is now supported.

When Partitioned database tables are used, data is no longer moved to the Staging database.

Partitioned database tables allow archiving of Production data without interference with Sterling Control Center normal operations.

Entire partitions (along with their indexes) are removed when the data has met its aging requirement, removing the requirement for periodic re-building of indexes.

Partitioned database tables are transparent to Sterling Control Center. The data on all active partitions is available to both reports and queries.

Sterling Control Center implements database partitioning in such a way that there is a separate partition for each day's data. To provide for monitoring servers whose clock is not set properly, Sterling Control Center ensures that there are partitions defined to hold data up to 30 days in the future from the Sterling Control Center clock. As data is collected from the monitored servers, the database automatically puts the data into the appropriate partition based on the date/time of the event.

When data is older than specified in the System Settings, entire partitions (along with their indexes) are removed, rather than deleting the data row by row.

A migration path is provided to keep the Staging data visible from a previous Sterling Control Center DB configuration when migrating to the use of partitioning.

Partitioning Overview

Sterling Control Center version 5.2 supports database table range partitioning. Table partitioning is a data organization scheme in which table data is divided across multiple storage objects called data partitions or ranges according to values in one or more table columns.

Partitioning improves performance and simplifies data management.

Advantages of partitioning:

- ◆ Improved manageability for large tables
- ◆ Increased query performance through data partition elimination. For example, if query includes date, don't look in partitions that don't include that date.
- ◆ Fast online data roll-in / roll-out

- ◆ Better optimization of storage costs
- ◆ Larger table capacity
- ◆ Each partition table will have fewer rows;
- ◆ Indexes on each partition table will be smaller (=faster seeks) than a corresponding index on the un-partitioned table.
- ◆ Greater scalability

Sterling Control Center version 5.2 supports database table range partitioning. It is a type of horizontal partitioning. Sterling Control Center supports partitioning for the following tables:

- ◆ EVENTS
- ◆ CD_STATS_LOG
- ◆ CE_STATS_LOG
- ◆ AUDIT_LOG (not partitioned in DB2 z/OS)
- ◆ EVENT_COMMENTS
- ◆ EVENTS_EXT

Sterling Control Center Supports database table partitioning for the following database server types (the support of partitioning may require additional features – check with your DBA):

- ◆ Oracle
- ◆ Microsoft SQL Server (MSSQL)
- ◆ IBM DB2 LUW (Linux, UNIX, and Windows)
- ◆ IBM DB2 z/OS
- ◆ MySQL

For MySQL database type, all the 5 tables are partitioned using PART_KEY column.

For all other database types, the partition column used is:

- ◆ EVENTS DATE_TIME column
- ◆ CD_STATS_LOG LOG_DATE_TIME column
- ◆ CE_STATS_LOG LOG_DATE_TIME column
- ◆ EVENT_COMMENTS LOG_DATE_TIME column
- ◆ AUDIT_LOG LOG_DATE_TIME column

The database table partitioning is supported only for Production database.

On an existing install, once the production database tables are converted to partitioned tables, the data movement from production to staging will be stopped. You should change your archiving process to archive of the Production database tables. The section “Archiving in Partition Environment” has more details on this.

For a new Sterling Control Center install, if you are planning to use Partitioned database tables on Production database, you do not need a staging database. When you are prompted for the Staging database connection details, you could specify 0.0.0.0 for the database hostname.

Partitioning is recommended for large database environment. Oracle recommends partitioning for tables greater than 2GB in size and for tables containing historic data.

System Requirements

The following are the requirements to database partitioning:

- ◆ Sterling Control Center version 5.2 or above
- ◆ Partitioning feature on the database server.
 - ◆ Check with your DBA to see if your database supports partitioning or not.
- ◆ Required database permissions
 - ◆ ALTER TABLE
 - ◆ DROP TABLE
 - ◆ CREATE INDEX
 - ◆ INSERT
 - ◆ UPDATE
 - ◆ DROP INDEX
 - ◆ ALTER INDEX
 - ◆ CREATE TABLE
- ◆ Database disk space – When partitioning is done, the existing tables are renamed and data from each of the existing tables is copied back to the partitioned table. After all the data from the old table is copied to the partitioned table, the old table is dropped. Before converting your Production database tables to partitioned tables, backup your existing production database.

Using Partitioning

To convert the Sterling Control Center production database tables as partitioned tables, the basic steps are:

- ◆ Install Sterling Control Center version 5.2 or above
- ◆ Backup your existing Production database (on existing Sterling Control Center installation).
- ◆ Run configCC.bat/sh script to configure the database
- ◆ Change the archive process.

Based on your environment, follow the steps of one of the following scenarios:

New Sterling Control Center Installation Using GUI Installer (Typically on Windows Platforms)

This scenario assumes that you will be partitioning your Production database tables. Since you will be using partitioning for Production database tables, there is no need for the Staging database.

Initial Setup

- ◆ Install Sterling Control Center using Windows Installer.
 - ◆ During the install, since there is no need for staging database, specify the following for the staging database connection details:
 - Database Host : 0.0.0.0
 - Port : 10000
 - User : x
 - Password: x
 - Database: x
 - ◆ After the install is done, run configCC.bat from the <<SCCInstallLocation>>/bin folder
 - ◆ During the DB Initialization step, you will be asked if you want to convert your Production database tables to partitioned tables, if partitioning is supported. Specify “Y” to convert the tables to partitioned tables. It may take some time to convert the tables to partitioned tables.

Ongoing Maintenance

Every day, based on the System Settings, older data from the five tables would be purged. If the old data is needed, then it must be archived before the purge process removes it.

New Sterling Control Center Installation Using Non-GUI Installer

This scenario assumes that you will be partitioning your Production database tables. Since you will be using partitioning for Production database tables, there is no need for the Staging database.

Initial Setup

- ◆ Install Sterling Control Center using the Installer.
- ◆ After the install is done, run configCC.sh from the <<SCCInstallLocation>>/bin folder
 - ◆ During the DB Initialization step, you will be asked if you want to convert your Production database tables to partitioned tables, if partitioning is supported. Specify “Y” to convert the tables to partitioned tables. It may take some time to convert the tables to partitioned tables.

- ◆ When configCC prompts for the Staging database connection details, specify the following for the staging database connection details:
 - Database Host : 0.0.0.0
 - Port : 10000
 - User : x
 - Password: x
 - Database: x

Ongoing Operation

Every day, based on the System Settings, the older data from the five tables would be purged. If the old data is needed, then it must be archived before the purge process removes it.

Existing Sterling Control Center Installation (With Data in Staging Database)

This scenario assumes that you will be converting your Production database tables to partitioned database tables. Once Production database tables are partitioned, there would not be any data movement from the Production database to the Staging database. You will still be able to access the data in the Staging database. Based on the System Settings, the data in the Staging database will be eventually purged.

Note: Before converting your Production database tables to partitioned tables, backup your existing production database data.

Initial Setup

- ◆ Install Sterling Control Center like any other maintenance install.
- ◆ After the install is done, run configCC.bat/sh from the <<SCCInstallLocation>>/bin folder
 - ◆ During the DB Initialization step, you will be asked if you want to convert your Production database tables to partitioned tables, if partitioning is supported. Specify “Y” to convert the tables to partitioned tables. It may take some time to convert the tables to partitioned tables.

Ongoing Operation

Every day, based on the System Settings, the older data from the five tables would be purged. If the old data is needed, then it must be archived before the purge process removes it. If your existing archive process was archiving out of the Staging database, then you must change it archive out of the Production database.

- ◆ No new data will be moved to the Staging database, and eventually the data in the Staging database will be purged based on age. When the data in the Staging database is no longer needed, you can run configCC.bat/sh and specify that the Staging database is not to be used by specifying the following for the staging database connection details:

- ◆ Database Host : 0.0.0.0
- ◆ Port : 10000
- ◆ User : x
- ◆ Password: x
- ◆ Database: x

Partitioning Process During Configuration

On a new Sterling Control Center install, when Production database is converted to partitioned tables, partitions for 30 days in the future are created. Thereafter, every day a new partition is created.

On an existing Sterling Control Center install, when Production database is converted to partitioned tables, the following is done:

Separate partition is created for each day's existing data

Partitions for 30 days in the future are created.

Partitioning Process During Engine Startup

Sterling Control Center does the following during engine startup:

- ◆ Remove older partitions based on the system settings. For example, if Sterling Control Center is configured to keep only 15 days of worth of data, then any partitions older than 15 days are dropped
- ◆ New partitions are created. Sterling Control Center ensures that always there are partitions for 30 days in the future from the current date.

Sterling Control Center also handles if the engine is down for more than a day. When the engine is started, it would create the necessary partitions for the days the engine was down.

Partitioning Process While Engine Is Running

When the engine is running, every day during the scheduled time, it does the following:

1. Creates new partition(s)
2. Drops older partition(s)

Archiving Data in the Partitioning Environment

This section provides high-level details on how the data from Production database can be retrieved in a partitioned environment for archival. In the partitioned database environment, it is recommended to archive everyday the previous day's data.

Overall, the process is to retrieve the old data from each table and then copy to the archive table. The following queries can be used to retrieve yesterday's data from each table:

```

SELECT * FROM EVENTS WHERE DATE_TIME >= '2009/09/04' AND DATE_TIME <
'2009/09/05';
SELECT * FROM CD_STATS_LOG WHERE LOG_DATE_TIME >= '2009/09/04' AND
LOG_DATE_TIME < '2009/09/05';
SELECT * FROM CE_STATS_LOG WHERE LOG_DATE_TIME >= '2009/09/04' AND
LOG_DATE_TIME < '2009/09/05';
SELECT * FROM EVENT_COMMENTS WHERE DATE_TIME >= '2009/09/04' AND DATE_TIME
< '2009/09/05';
SELECT * FROM AUDIT_LOG WHERE LOG_DATE_TIME >= '2009/09/04'
AND LOG_DATE_TIME < '2009/09/05';
SELECT * FROM EVENTS WHERE DATE_TIME >= '2009/09/04' AND DATE_TIME <
'2009/09/05';
SELECT * FROM CD_STATS_LOG WHERE LOG_DATE_TIME >= '2009/09/04' AND
LOG_DATE_TIME < '2009/09/05';
SELECT * FROM CE_STATS_LOG WHERE LOG_DATE_TIME >= '2009/09/04' AND
LOG_DATE_TIME < '2009/09/05';
SELECT * FROM EVENT_COMMENTS WHERE DATE_TIME >= '2009/09/04' AND DATE_TIME
< '2009/09/05';
SELECT * FROM AUDIT_LOG WHERE LOG_DATE_TIME >= '2009/09/04'
AND LOG_DATE_TIME < '2009/09/05';

```

In the example above, it assumes that current date is '2009/09/05' and the queries retrieve all data for yesterday (2009/09/04).

The above example is a generic approach to retrieve data regardless of database type. There are other alternate steps to retrieve the data for archival and it varies for each data type. Please refer the appropriate section below based on your database type.

Oracle Table Partitioning

In Oracle, Sterling Control Center names the table partitions using the following pattern:

For EVENTS table the partitions are named as EVyyyyymmdd where EV is the prefix, yyyy is the year, mm is the month and dd is the date of data. For CD_STATS_LOG table the partitions are named as CDyyyyymmdd where CD is the prefix, yyyy is the year, mm is the month and dd is the date of data.

For CE_STATS_LOG table the partitions are named as CEyyyyymmdd where CE is the prefix, yyyy is the year, mm is the month and dd is the date of data.

For EVENT_COMMENTS table the partitions are named as ECyyyymmdd where EC is the prefix, yyyy is the year, mm is the month and dd is the date of data.

For AUDIT_LOG table the partitions are named as ALyyyymmdd where AL is the prefix, yyyy is the year, mm is the month and dd is the date of data. If today's date is 2009/09/05, to retrieve yesterday's partition, use the following commands:

```
DROP TABLE TMP_ARC_EVENTS ;
ALTER TABLE EVENTS EXCHANGE PARTITION EV20090904 WITH TABLE
TMP_ARC_EVENTS ;
```

Then the data from the table TMP_ARC_EVENTS can be copied or move to the archive table.

Repeat above two commands for other tables.

Microsoft SQL Server (MSSQL) Partitioning

In MSSQL, a specific partition's data can be retrieved using partition function. Sterling Control Center uses different partition function for each table in MSSQL. EVENTS table uses EV_PART_FUNC, CD_STATS_LOG uses CD_PART_FUNC, CE_STATS_LOG table uses CE_PART_FUNC, EVENT_COMMENTS uses EC_PART_FUNC and AUDIT_LOG table uses AL_PART_FUNC function.

To retrieve and archive data from a specific partition on EVENTS table, use the following commands:

```
DROP TABLE TMP_ARC_EVENTS ;

ALTER TABLE EVENTS SWITCH PARTITION $PARTITION.EV_PART_FUNC
('2009/09/04') TO TMP_ARC_EVENTS
PARTITION$PARTITION.EV_PART_FUNC ('2009/09/04') ;
```

Then copy/move the data from TEMP_ARC_EVENTS table to archive table. The other tables can be archived in like fashion.

IBM DB2 – Linux, UNIX, Windows Partitioning

In DB2, Sterling Control Center names the table partitions using the following pattern:

For EVENTS table the partitions are named as EVyyyymmdd where EV is the prefix, yyyy is the year, mm is the month and dd is the date of data.

For CD_STATS_LOG table the partitions are named as CDyyyymmdd where CD is the prefix, yyyy is the year, mm is the month and dd is the date of data.

For CE_STATS_LOG table the partitions are named as CEyyyymmdd where CE is the prefix, yyyy is the year, mm is the month and dd is the date of data.

For EVENT_COMMENTS table the partitions are named as ECyyyymmdd where EC is the prefix, yyyy is the year, mm is the month and dd is the date of data.

For AUDIT_LOG table the partitions are named as ALyyyyymmdd where AL is the prefix, yyyy is the year, mm is the month and dd is the date of data.

If today's date is 2009/09/05, to retrieve yesterday's partition, use the following commands:

```
DROP TABLE TMP_ARC_EVENTS;  
ALTER TABLE EVENTS DETACH PARTITION EV20090904 INTO  
TMP_ARC_EVENTS;
```

Then the data from the table TMP_ARC_EVENTS can be copied or move to the archive table. The other tables can be archived in like fashion.

IBM DB2 z/OS Partitioning

In DB z/OS, the UNLOAD utility could be used to copy the data between tables.

MySQL Partitioning

In MySQL, the SELECT ... INTO OUTFILE SQL command could be used to dump specific range of data to a file and then LOAD DATA INFILE command could be used to load data into a archive table.

Additional Notes

Sterling Control Center uses the following table names during the processing and so you must avoid using these table names in your environment:

- ◆ EVENTS_OLD
- ◆ CD_STATS_LOG_OLD
- ◆ CE_STATS_LOG_OLD
- ◆ EVENT_COMMENTS_OLD
- ◆ AUDIT_LOG_OLD

Summary

Using Partitioned database tables allows Sterling Control Center to support larger monitoring environments, to collect more data efficiently, and to provide for 24x7 operations with no requirement for database maintenance.

References

<http://dev.mysql.com/tech-resources/articles/performance-partitioning.html>
<http://www.ibm.com/developerworks/data/library/techarticle/dm-0605ahuja2/index.html>
<http://www.mssqltips.com/tip.asp?tip=1796>
<http://msdn.microsoft.com/en-us/library/ms190787.aspx>
<http://publib.boulder.ibm.com/infocenter/db2luw/v9/topic/com.ibm.db2.udb.admin.doc/doc/c0021557.htm> <http://www.sqlteam.com/article/partitioning-the-data-in-a-table>
http://download.oracle.com/docs/cd/B28359_01/server.111/b28318/partconc.htm#CNCPT112