

**NEED TO KNOW**

# INTELLITRAX: AUTO ARCHIVE

## OVERVIEW

The DBat Auto Archive functionality allows the user to keep their database size manageable by archiving and optimizing older jobs. It features the following benefits:

- Makes regular and safe backups of customer data, which are then available for reporting.
- Keeps application performance optimal by keeping table size consistent.
- Allows customers to continue using the free Microsoft SQL Server engine without reaching its 2 GB database size limitation, and without having to upgrade to its more expensive edition.

## SETTING UP SCHEDULED AUTO ARCHIVE

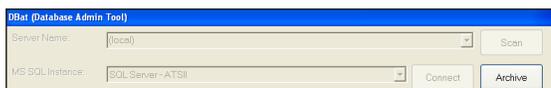
Scheduled auto archive uses SQL Server's SQLAgent functionality to run a script on a regular basis to optimize the database.

**IMPORTANT**

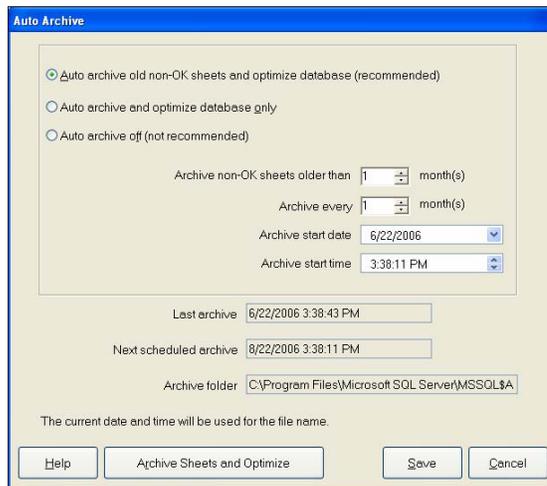
*Customers should turn on scheduled auto archive functionality as soon as the software is installed, and before any size or performance issues may occur.*

## PROCEDURE

1. Start DBat. Connect to a database and click **Archive**.



2. Click the **Auto archive old non-OK sheets and optimize database** radio button.



3. Set the auto archive options to match the customer's environment.

- Recommendation: keep the default values for “archive sheets older than X months” and “archive every Y months”.
- For the “archive start date and time,” choose a time when most downtime is available. Although the schedule auto archive can be done at the same time that users are running applications, users may experience slowdowns in this case.

4. Click **Save** to create the scheduled job and close the screen.

Advise customers to close all running applications during the Auto Archive scheduled time. To verify the process has finished, check the “Last Archive” field on the “Auto Archive” screen.

**Hot Tip**

*If the optimal start time for this job changes month to month based on customer work load, modify the schedule on a monthly basis to prevent any conflicts with running applications.*

# INTELLITRAX: AUTO ARCHIVE

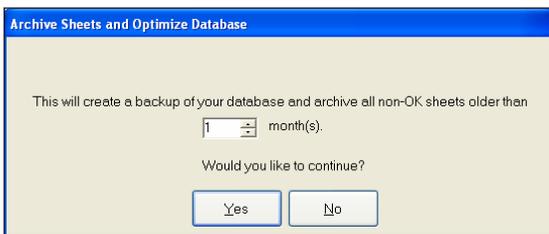
## MANUAL ARCHIVE

The archive script can also be triggered manually, and run along with the normal scheduled archive on an as-needed basis. It can also be used instead of the scheduled process, but this is not recommended.

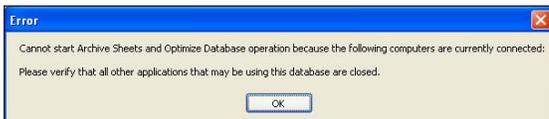
Although the scheduled and manual archive processes are essentially the same, they do optimize sheets in different ways. The scheduled script is meant for coexistence with running applications, and therefore uses more frequent smaller transactions to commit its changes. The manual script uses one large transaction, which is more efficient, but prevents any other applications from using or changing data. This forces it to also check that it is the only ATSH database user at the beginning of the process.

### PROCEDURE

1. From the Auto Archive screen, click the **Archive Sheets and Optimize** button.
2. Choose the optimize sheets month interval. The default setting is recommended.

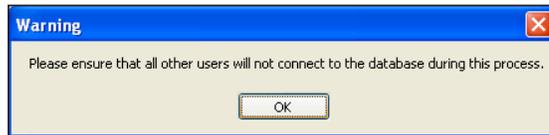


3. Make sure all users are disconnected from the database, and click **Yes**. If any users are still connected, an error such as this may appear:

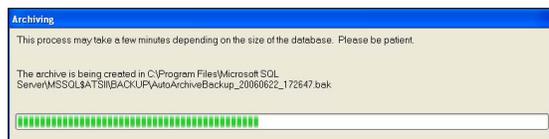


The easiest way to disconnect other users is to restart the SQL Server service and try this operation again.

4. Click **OK** on the prompt to acknowledge that no users should be connected to the database while the archive process is running.



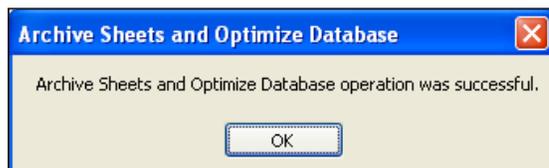
5. The following screen is shown while the archive process is running.



Be advised that the archive process may take a very long time (up to several hours), depending on the size of the database. Customers should not terminate the DBat application or use other applications during this time.

If the application is terminated prematurely, a long period may occur where the database cannot be used. This would happen because SQL Server would try to rollback the transaction and restore the database to the original state.

6. When this prompt appears, the archive process has finished successfully.



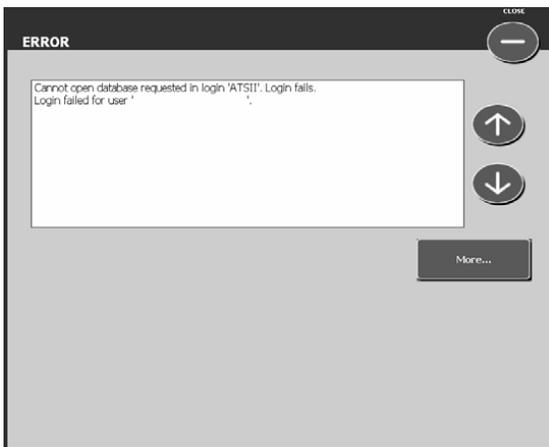
## MANUAL ARCHIVE TROUBLESHOOTING

If an error occurs during the archive process, e.g. if the server runs out of disk space, SQL Server might not rollback the transaction right away because it might think the operation is still pending. The best way to force the transaction to be rolled back is to restart the SQL Server service. After this point, it may take a long time for SQL Server to restore the database

# INTELLITRAX: AUTO ARCHIVE

back to a usable state. To verify the rollback process is happening, monitor the varying CPU and memory usage of the server in Task Manager.

To ensure no other users interfere with the archive process, DBat locks the database exclusively to itself. If the application is terminated before the process finishes, the database might not be usable for other applications even after the rollback process described above happens. This is because the first user to connect to the database will lock everyone else out. In Enterprise Manager, this database state will appear like this:



To unlock the database and stop errors, restart the SQL Server service to make sure that no one else is connected and connect to the database with DBat.

## RESTARTING SQL SERVER SERVICE

1. From Control Panel / Administrative tools, launch the “Services” applet.
2. Find the MSSQL\$ATSII service and click the **Restart** link.
3. Select **Yes** to “Restart Other Services”, which mentions the SQLAgent\$ATSII service dependency.

After successfully connecting to the database, the exclusive state is cleared and other applications should be able to connect.

## REPORTING ON ARCHIVED DATA

All backups, including backups created automatically by Auto Archive, are available in the ReportingTool for reporting as part of Change Source functionality. This is the recommended way for customers to access their archived data, and is safer than the process outlined in the upcoming “Restoring Archived Data” section.

Changing the data source restores the backup file into a separate database, named in a unique way based on the connected user’s computer. This allows multiple users to work with a different set of archived data from the same database server.

## PROCEDURE

1. Launch ReportingTool and click the **Change Source** button on the main screen.
2. The “Database Selection” screen lists each database backup available to be restored. Customers can optionally use the date and time in the automatic backup filename to help them determine which backup to restore.

In this example, the ATSII database is the primary database that all users are connected to. The first file below it represents a full backup made of the database on 6/20/2006 at 9:38 AM.

## Reference

*If the customer is having trouble determining the date when a particular job was run, they can reference the scan info dialog for the remaining OK sheet in the job, which never gets archived. Other production and make ready sheets should be present in the backup at around the same creation date.*

3. Select a database and click **OK**. The backup will be restored, which may take a long time.

**NEED TO KNOW**

## **INTELLITRAX: AUTO ARCHIVE**

If any database patches have become available since the backup was made, they are applied to the database after it is restored. This ensures that ReportingTool is working with the latest database structure.

4. When the main screen appears, the database selected for the reporting is now the restored database. Job and Sheet reports can be run as normal.

### **RESTORING ARCHIVED DATA**

The backups created by the Archive process go into the default SQL Server backup directory and become available in DBat for restoration.

If a customer wants to restore an older job to view it in PressTool (e.g. to view its make ready or production sheets), they should make a full backup of the current database and follow the normal DBat restore process.

Customers should not write any new data to this restored database. New data should only be written to the latest database.

When done using the archived data in PressTool, the customer should restore the backup made earlier.