

Maintain a Blue Prism database server

The instructions and recommendations contained within this section should be considered as a guide only. There are a large number of options available when maintaining databases that may relate not only to database specific settings, but also the host instance, as well as all instances on a given server.

It is recommended that both industry standard best practice and recommendations from an experienced database administrator are sought. These should be reviewed in light of the wider impact that they may have on the overall environment.

Customizing the database

When considering directly interacting with the database or its schema, it is essential that you are aware of the following limitations and advice.

Database modification

It is not supported to directly update databases created or provided for use with, or by, Blue Prism products. This includes adding, updating, modifying or deleting data; as well as making any changes to the database schema, including but not limited to adding new tables, changing table columns, changing indexes or foreign key constraints. Changes are only supported under instruction from Blue Prism Global Customer Support; or where provided in Blue Prism documentation and accompanied by a statement that endorses support when such changes are applied. The database schema, and the purpose and use of each column by Blue Prism is subject to change and should not be expected to remain consistent across Blue Prism versions.

Querying the database


It is not recommended to query the database manually or through use of third-party tools, however we recognize that there are some scenarios that may necessitate it; such as for reporting purposes. Where queries are run against the database(s), users are advised that it is their responsibility to:

- Review the complexity, cost, and frequency of queries to ensure that they do not impact or impede the performance or behavior of the system.
- Be aware of the queries and the frequencies and to proactively share the details with Global Customer Support when raising support tickets.
- Disable these queries if required as part of troubleshooting or when issues are being investigated.
- Permanently disable queries that are suspected of contributing or causing adverse behavior or performance.

SQL Server Maintenance

Database server settings

There are a number of SQL Server settings that should be reviewed to promote the performance and stability of the SQL Server platform. These settings vary in their scope and may affect: the Blue Prism database; the specific SQL instance that hosts it; and all SQL instances on a given server.


 It is strongly recommended that the respective Blue Prism documentation is reviewed in relation to the settings as they are critical in promoting a stable and reliable SQL platform.

The settings to be considered relate:

- Resource (RAM and CPU) allocation
- Storage utilisation and configuration
- Auto-growth
- Database statistics
- Use of Auto-Shrink

Shrink operations

Shrink commands, which may be featured as part of SQL maintenance plans, should be avoided because they cause significant amounts of fragmentation and are often unnecessary as databases re-use the released space as they re-grow. Additionally performance is impacted less during growth if the space has already been allocated and is available (rather than requiring dynamic expansion).

 It is recommended that activities which include shrinking either the database, data files, or log files, are avoided. There are a limited number of scenarios where it is appropriate to use a shrink operation, however shrinking typically results in performance degradation.

Backup

The Blue Prism databases must be backed up regularly and it is recommended that the backup storage location is a secure, remote location to cater for a range of disaster recovery scenarios.

The backup model should be appropriate to the criticality of the Blue Prism solution (which may vary across the environments) e.g. it is likely that the production environment will require a higher level of resilience and recovery than the development environment).

Blue Prism supports both **Simple** and **Full** SQL recovery modes and it is recommended that the benefits of each is reviewed to ensure the method chosen is appropriate to the criticality of the solution.

If the database has been set to use a **Full** recovery model, it is important that regular transaction log backups take place to periodically provide free space within the log files and restrict the amount of on-going log growth required.

General maintenance

The following maintenance tasks should be carried out periodically – these can typically be achieved using standard IT tools:

- **Available physical disk space** – The amount of physical disk space available on each disk array or partition used by the SQL Server should be reviewed to ensure appropriate free space is available. This is especially important where a disk partition is shared by multiple databases or applications.
- **Disk array fragmentation** – The amount of disk fragmentation should be reviewed to establish if a defragmentation task is required. Fragmentation causes performance degradation in read/write operations to the disk which in turn will impact the general performance of the SQL server. This is of particular importance for the disk array on which the SQL log files reside.
- **Paging file** – The paging file is used by the operating system as part of the management of RAM and therefore it should be set to an appropriate size
- **Event viewer** – The event viewer on the SQL Server should be reviewed for errors as these could be indicative of underlying issues.
- **Verify maintenance plan and backup execution** – It is important to make sure that the scheduled maintenance and backup tasks have taken place as expected, and without error. Database backups should also be verified to ensure that they can actually be restored if required. Options such as RESTORE VERIFYONLY and RESTORE can be considered for this type of verification although both are intensive processes so may be best applied within a non-production environment.

Database maintenance

In order to maximise the performance of the database it is important that regular maintenance takes place. Maintenance tasks can also be used to identify issues that may be occurring at the database level before they have become apparent.

The following maintenance activities should be undertaken regularly, ideally on a scheduled basis.

Verify data integrity

Use of DBCC CHECKDB will check the logical and physical integrity of all of the objects in the specified database. If any issues are found there are a number of repair procedures available, however it should be noted that the resolution for some issues may require the database to be overwritten with a restore of a previously taken backup.

Identify and resolve index fragmentation

The performance of the indexes is important to the overall performance of the database and therefore it is recommended that the level of fragmentation is reviewed and appropriate steps taken to defragment the indexes as required.

Whilst there are few out of the box features that provide insight into the level of index fragmentation, online resources contain examples of scripts that can help with checking the level of fragmentation across all indexes in a database.

Commands such as ALTER INDEX can be used to rebuild or re-organise the indexes.



Consider also the fragmentation of the underlying file system as this may further impact the performance of the SQL Server.

Ensure statistics are up-to-date

Database Statistics are used by the Query Processor in Microsoft SQL Server to determine how a query should be executed. It is therefore imperative that these are always up-to-date and accurate.

The most common way to generally ensure that statistics are maintained is to turn on the following database settings:

- Auto Create Statistics
- Auto Update Statistics

When updating statistics manually it is important that the action includes carrying out a full scan – this can be ensured by using the alter index action to carry out an index rebuild as this uses a full scan and updates the index statistics automatically. This action however does not update column statistics – this can be done using the UPDATE STATISTICS WITH FULLSCAN, COLUMNS statement.

Free space within SQL Data/Log files

The amount of free space available within the data and log files should be regularly evaluated to establish whether the size of the data or log files needs to be increased to provide additional free space. The act of expanding the files dynamically at the time-of-need can have a performance impact, particularly when log files grow, so it is recommended that where possible, growth should take place manually and at a predetermined time.

Evaluating file size and % free space

Log Files

The command `DBCC SQLPERF(logspace)` can be executed against the master database to provide the required information for all log files in a given SQL instance.

Data and Log Files

The method `FILEPROPERTY` can be used as part of a query executed against the master database to provide the required information for all data and log files in a given SQL instance.

```
SELECT DB_NAME() AS DatabaseName,  
name AS FileName,  
size/128.0 AS FileSize_MB,  
size/128.0 - CAST(FILEPROPERTY(name, 'SpaceUsed') AS INT)/128.0 AS FreeSpace_MB  
FROM sys.database_files;
```

Whilst databases should be allowed to grow automatically if they need to, the ideal scenario is that the file sizes are increased manually in advance ready for use by SQL when the space is required.

Periodic review of the SQL database settings

The SQL Database Settings referenced within this document should be reviewed periodically to ensure that no settings have inadvertently been changed, and also as an opportunity to ensure that any additional SQL databases or instances have not been configured with settings that may conflict with the instance hosting the Blue Prism database(s).

Blue Prism in-product maintenance

The Blue Prism software includes a number of maintenance capabilities which are predominately focused on removing old information from the database to prevent ongoing database growth and expansion.


Archiving

Blue Prism Archiving is used to extract specific data types of a certain age from the database for the purposes of retaining historic data externally for prolonged periods of time without requiring ever-increasing amounts of space within the SQL database.

The facility to re-import archived data is provided which allows the Blue Prism GUI to be used to view and traverse the historic information.

The following data is currently affected by archiving:

- Session logs are extracted into a compressed XML file and removed from the database.
- Process alerts that relate to the sessions logs being extracted are permanently deleted.

 The process of archiving relies on the SQL database and therefore manual archiving should be performed under controlled conditions to limit the impact of this action on other users. Additionally the SQL actions performed by archiving will cause SQL transaction logs to be generated and therefore it is recommended that, if the database is set to Full Recovery Mode, that a log backup is taken shortly after large archiving operations have completed to reduce the log size within the file.

Blue Prism recommends that a centralised archive location is defined and configured for appropriate access by the relevant users and resources that will be responsible for archiving. The permissions granted for the location will need to include the ability to create folders, and read/write files.

There are three options for archiving:

- Manually via the user interface.
- Automatically via the user interface.
- Using the command line (automatec.exe).

The automatic and command line archiving actions provide the advanced option to permanently delete the applicable data without first extracting the data to folder location.

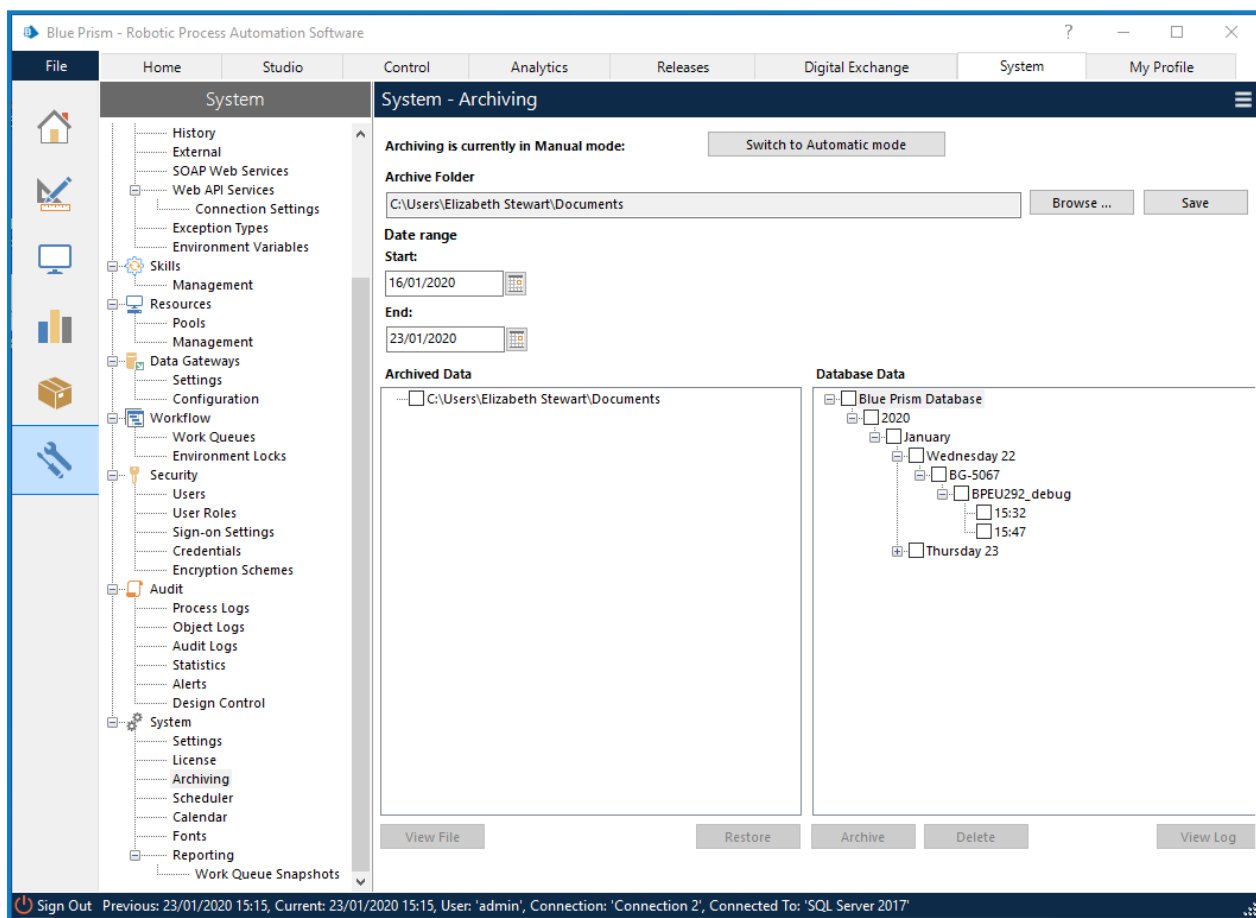
Manual initiation

Within Blue Prism the logs are displayed using a tree structure based on the:

- year, month and day that a given process was started.
- name of the process.
- runtime resource which executed it.
- time at which it started.

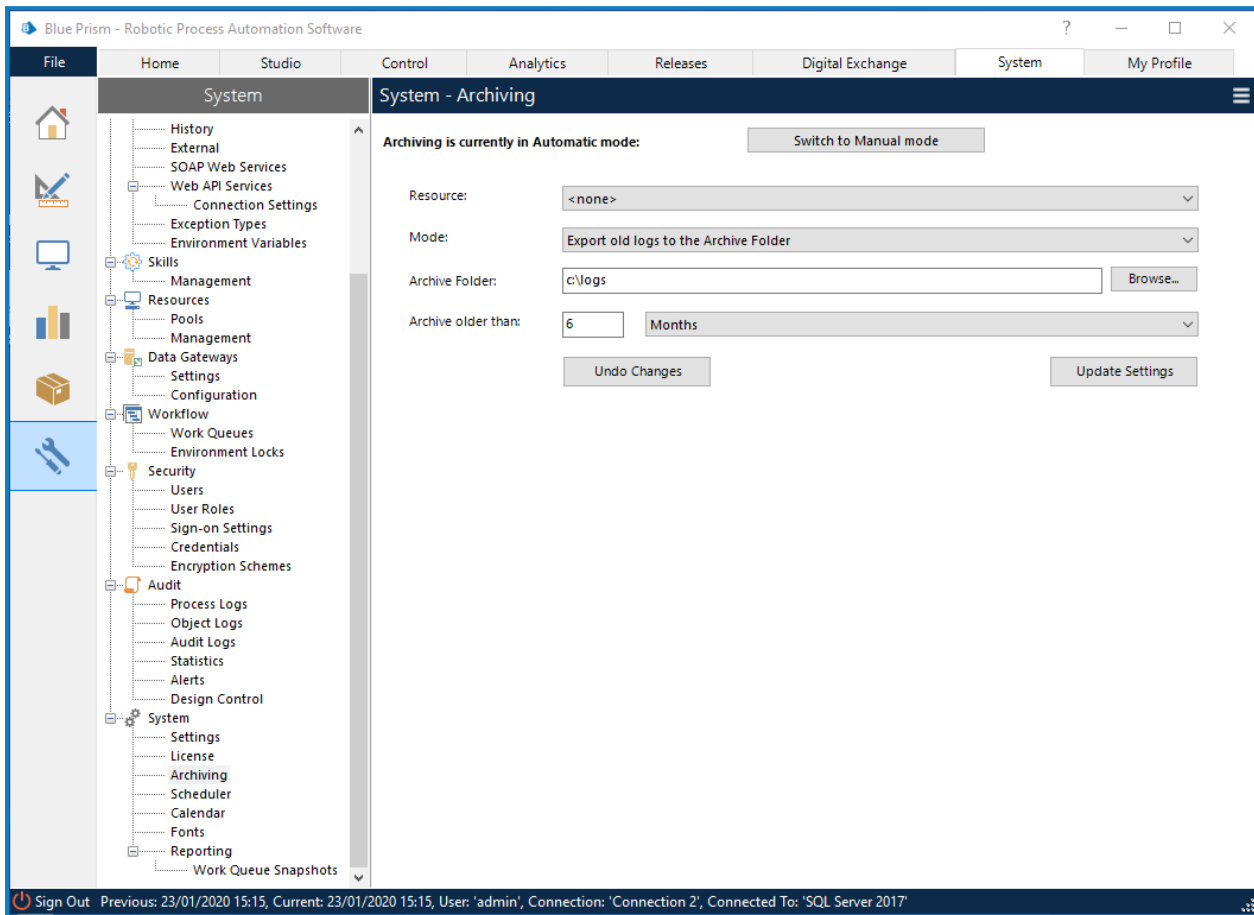
Manual archiving is initiated by navigating to the **Archiving** tab (System Manager -> System -> Archiving) and ensuring that **manual** mode is selected.

The user can select a number of individual items or branches in the tree for export, as well as the destination folder where the archived information will be stored. Blue Prism creates a folder structure within the target location representing a similar structure to how the data is presented via the Blue Prism user interface.



Automatic, scheduled initiation

Scheduled archiving is configured by navigating to the **Archiving** tab (System Manager -> System -> Archiving) and ensuring that **automatic** mode is selected.



With automatic mode there are a number of configurable options:

- **Resource** – The name of the resource that will be used for the archiving operation.
- **Mode** – Provides the option to either: export and archive the files prior to automatically removing the data from the system; or alternatively to permanently delete the data without archiving.
- **Archive folder** – The location where the archived data will be placed. Blue Prism creates a folder structure within the target location representing a similar structure to how the data is presented via the Blue Prism user interface.
- **Archive older than** – A configurable value that indicates what age the data must be to be processed by the scheduled archive action.



The frequency at which archiving procedures are initiated is automatically determined by Blue Prism to encourage frequent, low-volume, archiving rather than occasional processing of larger amounts of data.

Command line initiation

A command line method of AutomateC.exe can be used to trigger an archive. Archived files exported using the command line will be placed in the folder location specified on the **Archiving** tab within Blue Prism **System Manager**.

If configured for **single sign-on** the command to archive data older than 3 months is:

```
AutomateC.exe /archive /sso /age 3m
```


If configured for **native** authentication the same command is:

```
AutomateC.exe /archive /user "username" "password" /age 3m
```

When triggered by the command line the archiving process is carried out by the resource on which the command is executed. It is therefore necessary for the chosen resource to have full read/write/execute access to the folder location specified within Blue Prism.

If there is a desire to simply **delete the archived information rather than storing it to a folder location**, the /delete switch can be used:

```
>AutomateC.exe /archive /sso /age 3m /delete  
AutomateC.exe /archive /user "username" "password" /age 3m /delete
```

 Use of the /delete switch permanently deletes the data and therefore removes the ability to retain or re-import the archived information at a later time.

Audit logs

There is currently no in-product mechanism for deleting audit logs. The tables are generally quite small and grow slowly, except for the inclusion of historic versions of process and business objects.

Database updates

Periodically Blue Prism software updates also require updates or changes to the Blue Prism database. Such database updates can be manually driven from within the product, or alternatively SQL scripts can be provided to allow the updates to be applied manually on the SQL Server.

It is recommended that a full database backup takes place and is subsequently verified prior to any manual database updates or modifications.

Session logs


There are a number of options for managing and deleting session log information:

- **Using the Archiving feature (recommended)** – This includes manual, automatic and command line archiving (or permanent deletion) of session logs and can be aligned with the organisation's data retention policy.
- **Via the User Interface** – Within Control Room, select the items to be deleted and choose to Delete.

Work queues

Historic work items can be permanently deleted using any of the following methods:

- **Via the User interface** – Within Control Room, select the items to be deleted and choose to Delete Worked Cases.
- **Using a Blue Prism process** – A custom Blue Prism process can be configured to review and delete the appropriate items using the Delete Processed Items action.
- **Command line** – AutomateC.exe includes a /queueclearworked command which can be used to permanently delete worked items.

 It is important to ensure that work items are not required as part of a data retention policy prior to deleting them as there is presently no facility for archiving or re-importing them.