

Sterling B2B Integrator



Archive and Purge

Version 5.2

Sterling B2B Integrator



Archive and Purge

Version 5.2

Note

Before using this information and the product it supports, read the information in "Notices" on page 15.

Copyright

This edition applies to Version 5 Release 2 of Sterling B2B Integrator and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2000, 2015.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Archive and Purge	1
Archive Business Process Data	1
Before You Archive	2
Configure Archive Schedules	3
Archive Business Process Data Manually	6
Configure the Archive Log.	6
View the Archive Log	6
Archive Reporting	7
Restore Archived Business Process Data	7
View Restored Business Process Data	8
Remove Restored Business Process Data	8
Troubleshooting	9

External Purge (V5.2.4 or later)	10
Starting External Purge	10
Monitoring External Purge	11
Stopping External Purge	11
Disabling External Purge	11
Tuning External Purge.	12
bin Directory Files Used by External Purge.	13

Notices	15
Trademarks	17
Terms and conditions for product documentation.	18

Archive and Purge

Archiving your business process data protects your critical data, conserves database disk space, and file system disk space (when using documents on disk) thereby improving the efficiency of IBM® Sterling B2B Integrator.

Archive and purge cleans up business process data. The archive and purge operation searches for eligible records to archive from the live database tables. You can schedule archive and purge operation or perform them manually.

Archive or purge is not the only source of cleanup. Other services may need additional requirements or considerations and are discussed in the respective service documentation. For example, to purge mailbox messages, you should use the Mailbox Delete service, which deletes messages in one or more mailboxes.

Archive Business Process Data

Sterling B2B Integrator uses a set of components to move business process data that has reached its specified lifespan out of the live database.

These components include:

- Index Business Process service
- Backup Business Process service
- Purge service

Data that has been removed using the Backup Business Process component can later be restored to the same version using the Restore Business Process component.

Business process data written to the database is determined by the persistence level specified in the Business Process Manager or in the noapp.properties file. Decreasing the persistence level increases the business process performance at the cost of full-tracking for each step of the business process.

Business processes are eligible for archiving as soon as their archive flag is set by the Index Business Process service. You do not have to wait until a business process exceeds its lifespan before archiving it. After the archive flag is set, the business process is then archived the next time the Backup Business Process service is run either by schedule (at 2 A.M. each morning) or manually. The data is still available in the system because it cannot be purged (with the Purge service) until it has exceeded its lifespan.

The Restore Business Process service restores archived business process data from physical media to a restored data location, where it can be searched and viewed.

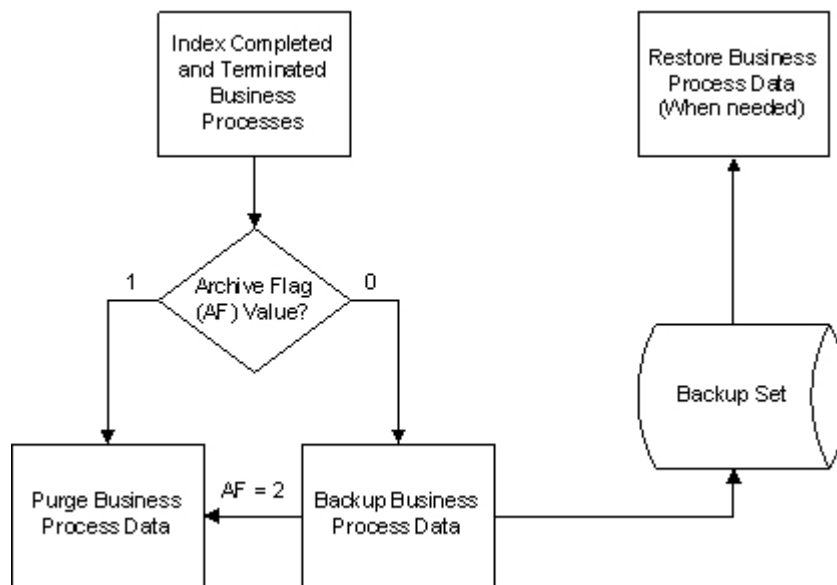
The Index Business Process service runs a process that looks for archiving parameters defined in the business process. When the process finds these parameters, the process creates summary information and writes the information to a table in the database.

The Backup Business Process service runs based on how you have configured your archive settings. The Backup Business Process service retains the business process

and its related data to a file system as Java™ serialized objects for backup or purging. The Purge service then runs and removes data from the database, file system, or both as specified in the archive settings.

You cannot archive active business process instances (for example, halting, halted or interrupted). Only data for completed or terminated business processes that have been indexed can be archived.

The following figure shows the archiving process.



The Archive Flag is initially set by the Index Business Process service to 0 or 1. Valid values are:

- 0 - Back up data
- 1 - Purge data without performing a back up
- 2 - Data back up sets the flag to 2, which purges data

You cannot delete a workflow definition of a business process (using the Source Manager of the Business Process Manager) while an instance of that process is being indexed, archived, or purged.

In the same manner that you schedule business processes, you can schedule the intervals at which you want to index, archive, and purge the contents of the database.

You can also define the lifespan, or the length of time, in days and hours, to retain the data in the live database tables. This is the length of time until the business process instance expires.

Before You Archive

You can configure to automatically archive business process data as often as you need, or you can archive business process data manually.

For more information about changing the default settings in your archive schedule, see “Configure Archive Schedules” on page 3.

For more information about archiving business process data manually, see “Archive Business Process Data Manually” on page 6.

You should not create your own service to back up the data (either manually or on a schedule). Similarly, you should not create your own service and then manually archive data at the same time that an archive process is automatically running (from a schedule in the Archive Manager). It can corrupt archived information as each process could be trying to archive the same data. Use the lock feature in the Archive Manager to prevent simultaneous manual and automatic archiving when using standard backup tools.

Before you begin to archive business process data, you should consider the following:

- Indicate default values that apply to all business process instances that you run.
- Schedule archive process depending on the volume of business processes that you run.
- Archive business level information for audit purposes, including:
 - Documents that support non-repudiation for AS1 and AS2 protocols
 - Correlations that provide searching ability
- Maintain data in live database tables, to enable faster retrieval for data that you want to view in the immediate future.
- Separate JVM is not required to back up data, thereby allowing you to allocate more resources to .
- Archive streamed-in documents that are much larger than the amount of available memory, because documents are not reassembled in memory prior to writing them on the disk. Instead, the backup process reads and writes raw DATA_TABLE pages.
- In a clustered environment, run the backup process on any one node. If that node goes down, run it on a different node.

If you change the name of the directory used to store document payloads when document storage is set to file system (which is stored in the document_dir property of the jdbc.properties file), you must make the same change in the following properties:

- RESTORE_DOCUMENT_DIR property of the jdbc.properties file
- PURGE_DOCDISK_LIST_FILENAME property in the archivethreadproperties file

Configure Archive Schedules

You can define the location of the archive data when configuring archive schedule settings. After you configure and save the archive schedule, the archive process runs automatically based on the parameters you provided.

About this task

You can still archive manually, however, you should consider the scheduled archive before running the archiving activities manually.

For more information about running archiving manually, see “Archive Business Process Data Manually” on page 6.

To configure archive settings (either the default settings or your manual settings):

Procedure

1. From the **Administration** menu, select **Operations > Archive Manager**.
2. In the Archive Manager page that is displayed, under **Archive Configuration**, next to **Configure Archive Settings**, click **Go!** This message is displayed:
 UI lock enabled. Lock will be released only by clicking Cancel or Finish in the wizard.
3. Click **OK**.
 If the lock is not released or if you close the page without clicking **Cancel** or **Finish**, you must use the Lock Manager function to release the lock.
4. In the Configure page that is displayed, enter information in the following fields and click **Next**.

Field	Description
Default values that business processes should remain in the system:	It indicates the lifespan of the business process to remain active in the system. Enter the values in Days and Hours fields.
Default values that trackable business process information remains in the system:	<p>It indicates the additional time of trackable business process data, such as security information certificates and non-repudiation supporting documents, to remain active in the system.</p> <p>The trackable business process data includes archive or purge-related tables in a business process and is helpful when performing the data searches. These tables are retained longer than the tables in the same process that are deleted using the default lifespan set, for example, if you enter two days for lifespan of the process and one day for the trackable business process data, then trackable business process data remains active in the system for three days.</p> <p>Enter the values in Days and Hours fields.</p>
Directory where archived data will be stored:	It indicates the directory where the archived data is stored. Enter the complete path to the archive location in the Backup Directory field.
Expired Business Processes should be:	It indicates if the expired business process should be archived or purged. Select the Archived radio button to archive the expired business process or the Purged radio button to purge the archived business process.

5. In the Business Process Definitions page that is displayed, the left-side window displays the list of business processes that do not have the lifespan set as system default and the right-side window is populated with business process that have the lifespan set as system default.
 To apply your archive settings to a business process, select and move it to the right-hand window. If you choose to remove archive settings for a business process, move it to the left-side window. You can also filter and search for business processes by name in the left-side window.
 The lifespan setting from the Configure page applies to all business processes that are set with a system default lifespan. Any processes in the right-side window will be set to the system default lifespan, and will pull the lifespan setting at indexing that is set in the Configure page. If you change the removal method (archive or purge), business processes that are not moved to the right-side window retain their previous removal method.
6. Click **Next**. The business process update message is displayed:

Updating of business processes will occur in the background and may take 30 minutes or more, depending on the number of processes selected.

7. Click **OK**.
8. In the Schedule Type page that is displayed, select the checkbox next to **Use 24 Hour Clock Display** if you want schedule archive in a 24 hour format.
9. Select the archive process schedule by choosing one of the following radio buttons:
 - **Run based on timer** - To archive or purge business process data based on the timer set (in hours and minutes).
 - **Run daily** - To archive or purge business process data daily.
 - **Run based on day(s) of the week** - To archive or purge business process data on specific days in a week.
 - **Run based on day(s) of the month** - To archive or purge business process data on specific days in a month
10. Click **Next**.
11. In the Schedule Settings page that is displayed, the fields are displayed depending on the schedule type selected in Schedule Type page.

For all schedule types except **Run based on timer** type, you can schedule archive activity at regular intervals. To do this, select the checkbox next to **Check here to select time interval** and specify the start time, end time, and the interval.

To run the scheduled archive process when starts, as well as on the specified schedule, select the checkbox next to **At startup**. This option is not available if you select **Run based on day(s) of the month** on the Schedule Type page.

 - If you selected **Run based on timer** in the Schedule Type page, enter numeric data in the **Hour(s)** and **Min(s)** fields to set how often the process should run.
 - If you selected **Run daily** in the Schedule Type page, enter when you want to run the process each day. Click **add** to specify an entry. You can enter more than one item. Click **delete** to remove an entry.
 - If you selected **Run based on day(s) of the week** in the Schedule Type page, enter when you want to run the archive process each day. For each time entry, select the day from the **Select Day(s)** drop-down list when you want to run the process. Click **add** to specify an entry. You can enter more than one item. Click **delete** to remove an entry.
 - If you selected **Run based on day(s) of the month** on the Schedule Type page, enter the time when you want to run the process each month. For each time entry, select the day of the month from the **Select Day(s)** drop-down list when you want to run the process. Select **LDOM** to choose the last day of the month. Click **add** to specify an entry. You can enter more than one item. Click **delete** to remove an entry.
12. Click **Next**.
13. In the Schedule Exclusions (displayed if you selected **Run based on timer** in the Schedule Type page) and Date Exclusions page that is displayed, depending on your selection in the Schedule Type page, select the time and date when the archive process should not run.

Click **add** to specify an entry. You can enter more than one item. Click **delete** to remove an entry.
14. In the Confirm page that is displayed, review the changes you made, and perform one of the following tasks:

- To change your settings, click **Back**.
- To abandon the changes you carried out, click **Cancel**.
- To save the changes you carried out, click **Finish**.

Archive Business Process Data Manually

After you have configured archive settings, you can archive completed or terminated business process instances at any time.

About this task

To archive business process instances:

Procedure

1. From the **Administration Menu**, select **Operations > Archive Manager**.
2. In the Archive Manager page that is displayed, under **Archive Execution**, next to **Archive Business Processes**, click **Go!** The archive process confirmation message is displayed:
Are you sure you want to execute Archive process now?
Click 'OK' to proceed or 'Cancel' to abort.
3. Click **OK** to run the archive process or click **Cancel** to abandon the archive process.

Configure the Archive Log

To compile an archive log of business process data, configure the archive log.

About this task

The archive log provides details about business process instances, including errors encountered during archiving.

To configure the archive log settings:

Procedure

1. From the **Administration Menu**, select **Operations > System > Logs**.
2. On the System Logs page, under **Application Logs**, click the icon next to **Archive**. The Archive: Log Settings pop-up window is displayed.
3. On the Archive: Log Settings page, perform the following actions:
 - a. Select the maximum number of lines the log can contain from the **Roll Over Interval** drop-down list. After attaining the maximum number of lines set, it rolls over to a new log.
 - b. Select the **On** radio button to enable logging. Logging is disabled if you select **Off** radio button.
 - c. Click **Save**. The Environment pop-up window closes and you are redirected to the user interface.

View the Archive Log

The archive log provides details about business process instances, including errors encountered during archiving.

About this task

It compiles an archive log of business process data archived every day.

To view the archive.log file, use one of the following methods:

Procedure

- To view the archive log using the Archive Manager:
 1. From the **Administration Menu**, select **Operations > Archive Manager**.
 2. On the Archive Manager page, under **Archive Log**, next to **View Archive Log**, click **Go!**

The Archive Log pop-up window displays the archive.log information.

- To view the archive log using the System Logs:
 1. From the **Administration** menu, select **Operations > System > Logs**.
 2. On the System Logs page, under **Application Logs**, click **archive.log**.

The Archive Log pop-up window displays the archive.log information.

Archive Reporting

The Archive Status Report describes the values of all of the configuration parameters and the results of the archive for each table.

The Archive Status report has two sections:

- The values of all of the configuration parameters, along with the results of the archive run and the name of the backup set directory.
- The results of the archive for each table, including the name of the table, the expected and actual row counts, and the elapsed time for each table. The tables are sorted by the actual table names. If no data is eligible for archiving, the second section does not contain any table information.

The status report does not show error and exception information, which is shown in the archive log.

Status reporting uses the standard application programming interface to capture stdout and stderr to report results. Reporting does not capture results from the Command Line Adapter 2.

Restore Archived Business Process Data

You can restore archived business process data.

About this task

To restore archived business process data:

Procedure

1. From the **Administration Menu**, select **Operations > Archive Manager**.
2. On the Archive Manager page, under **Restore Manager**, next to **Manage Restoring Business Processes**, click **Go!**
3. On the Restore Manager page, enter the full path for the restore directory in the **Restore Directory** field. The restore directory stores archived data that you want to restore.

4. Click **Go!** to restore the data.

View Restored Business Process Data

You can view and search for specific restored business process data.

About this task

To view archived business process data restored in :

Procedure

1. From the **Administration Menu**, select **Business Processes > Monitor > Advanced Search > Business Processes**.
2. Under Search Location, select **Restore Tables**.
3. In the Business Process Monitor Advanced Search page, under **Search Location**, next to **Select the area to search from**, select the **Restore Tables** radio button.
4. In the Business Process Monitor Search (Restore) page, complete one of the following steps:
 - Click **Go!** at the bottom of the page to view all restored business process data.
 - Search for specific restored business process data. Indicate any combination of the following criteria and click **Go!**:
 - Process ID – Display data identified by the number assigned when the business process ran.
 - Business Processes – Display data for business processes with names that contain the specified character or string.
 - System Business Processes – Display data for system business processes with names that contain the specified character or string.
 - State – Display data for business processes with a final state (that is, All, Completed, and Terminated) that corresponds to the specified state.
 - Status – Display data for business processes with a final status (that is, All, Success, and Error) that corresponds to the specified status.
 - Start date and time range – Display data for business processes run within the specified start dates and times.

Remove Restored Business Process Data

You can remove restored business process data.

About this task

To remove business process data restored from archive tables:

Procedure

1. From the **Administration Menu**, select **Operations > Archive Manager**.
2. In the Archive Manager page, under **Restore Manager**, next to **Manage Restoring Business Processes**, click **Go!**
3. In the Restore Manager page, under **Cleanup**, next to **Clean restore** area, click **Go!**
4. When prompted with the message: Are you sure you want to clean the restore area?, click **OK**.

Results

After cleaning the restore area, the following confirmation message is displayed:

Successfully completed cleaning the restore area.

Troubleshooting

After you have addressed issues described in Before You Archive and the archiving or purging process does not appear to work properly, check for these following issues.

If you have addressed the issues described in “Before You Archive” on page 2 and the archiving or purging process did not appear to work properly, check for the following issues:

- From the **Administration** menu, select **Business Processes > Monitor > Central Search**.
 - In the Central Search page, under Search Options, next to Business Process box, enter and click **Go!**
 - In the Central Search Results page, click the business process that matches your search. Verify that Schedule_BackupService business process completed successfully. Schedule_BackupService must run for data to be removed from the live database tables.
- From the **Administration Menu**, select **Deployment > Schedules**.
 - In the Schedules page, under Search, next to **by Name**, enter and click **Go!**
 - Verify that the default time for this service to run is every ten minutes.
- From the **Administration** menu, select **Business Processes > Monitor > Central Search**.
 - In the Central Search page, under **Search Options**, next to the **Business Process** box, enter and click **Go!**
 - In the Central Search Results page, click the business process that matches your search. Verify that Schedule_IndexBusinessProcess business process completed successfully. Schedule_IndexBusinessProcess must run to set the expiry dates for completed business processes.
- From the **Administration Menu**, select **Business Processes > Monitor > Advanced Search > Business Processes**.
 - In the Business Process Monitor Advanced Search page, under **Search Location** , next to **Select the area to search from**, select **Live Tables**.
 - Perform the following tasks under **Search Using Business Process Name**:
 1. Select **ALL** for Business Processes, System Business Processes, State, and Status.
 2. Set the Start Date and End Date to include business processes that should have been archived according to the default value set in the Archive Manager configuration.
 3. Click **Go!**
 4. In the Monitor page, under the **Expires** column, click **Info** for the business process archive schedule information.
 5. In the Business Process Expiration pop-up window that is displayed, verify that the business process is scheduled for archive. If the information is not displayed or if it reflects a past date, the business process has not

been configured for archiving. You must terminate the business process manually, configure the archive settings, and then run the business process again.

- Check the archive log for errors.

External Purge (V5.2.4 or later)

In addition to the default purge process, which is automated, you can run an external purge process manually from a command line.

Important: IBM recommends that you not run External Purge and the default purge at the same time, because of risks to your data integrity.

External Purge includes the following features:

- Uses a configurable number of threads to allow for scaling.
- Runs in a separate JVM from the engine and adapters.
- Can run when the ASI nodes are down if the database is available.

Starting External Purge

You can start an external purge by running a command.

Procedure

1. If you have a cluster installation, make sure that you are not running External Purge on another node of the cluster. If you are running External Purge on more than one node, it will cause database locks. For more information, see “Monitoring External Purge” on page 11.

Important: It is the customer's responsibility to keep exactly one of these external purge processes running in the cluster. If you need to shut down the hardware that is running External Purge, you must run it on a different node.

2. Disable the default Purge service. Click **Deployment > Schedules**, search for PurgeService, and then clear the **Enabled** checkbox for PurgeService.

Note: If you have customized the Schedule_PurgeService business process, you need to add your customized code to the new Schedule_PurgeService business process. The new Schedule_PurgeService business process includes a warning message about running External Purge while the scheduled new purge is enabled.

3. Turn off the "Important Schedules are Disabled" email notifications by adding the following line to the customer_overrides.properties file. This removes the Purge Service SQL query for the schedule.

```
resourceMonitor.ScheduleMonitor.propertyValue.1=SELECT STATUS,  
SERVICENAME FROM SCHEDULE WHERE SERVICENAME IN  
( 'BackupService', 'IndexBusinessProcessService', 'AssociateBPsToDocs',  
, 'BPrecovery', 'BPLinkagePurgeService' )
```

4. Do one of the following:
 - In Windows, do one of the following:
 - Run the command `startExternalPurgeWindowsService.cmd` from the `INSTALL_DIR\bin` directory.
 - Run the Windows service. Click **Control Panel > Administrative Tools > Services**, right-click **IBM Sterling B2B Integrator External Purge at BASE_PORT**, and click **Start**.

- In UNIX, run the command `control_extpurge.sh` start from the `INSTALL_DIR/bin` directory.
- In iSeries, run the command
`SBMJOB CMD (QSH CMD('cd INSTALL_DIR/bin; ./ExternalPurgeConsole.sh -clearlocks > INSTALL_DIR/logs/extpurge.log')) JOB(EXTPURGE).`

Monitoring External Purge

You can monitor the external purge through the **Database Usage, Lock Manager,** and log file in the `INSTALL` logs.

After you start or stop External Purge, you can monitor its operation using the following tools:

- Log files in `INSTALL/logs` (`StopExternalPurge.sh.log`) and `INSTALL/logs/PURGE` (`extpurge.log.*`).
- The **Database Usage** report.
 1. Click **Operations > System > Troubleshooter**.
 2. Click the **Database Usage** link on the **System Troubleshooting** page.
- The **Lock Manager** page.
 1. Click **Operations > Lock Manager**.
 2. Search for the purge process.

Stopping External Purge

You can stop an external purge and return to the default purge.

Procedure

1. Do one of the following:
 - In Windows, do one of the following:
 - Run the command `stopExternalPurgeWindowsService.cmd` from the `INSTALL_DIR\bin` directory.
 - Run the Windows service. Click **Control Panel > Administrative Tools > Services**, right-click **IBM Sterling B2B Integrator External Purge at BASE_PORT**, and click **Stop**.
 - In UNIX, run the command `control_extpurge.sh stop` from the `INSTALL_DIR/bin` directory.
 - In iSeries, use the `WRKACTJOB` command to end the `EXTPURGE` job that you used to start the external purge.
2. If you are disabling External Purge and returning to the default purge, see “Disabling External Purge.”

Disabling External Purge

You can disable external purge by removing the locks that are created by the external purge and enable the default purge process.

Procedure

1. Stop External Purge. For more information, see “Stopping External Purge.”
2. Remove the locks that were created by the external purge. Click **Operations > System > Lock Manager** and remove the following locks:
 - `hpp.Purge`
 - `PURGE_SERVICE`

3. Re-enable the default purge process. Click **Deployment > Schedules** and select the **Enabled** check box for PurgeService.

Tuning External Purge

You can tune External Purge using the `archivethread.properties_hpp_ext` file.

Procedure

1. Open the `archivethread.properties_hpp_ext` file in the `INSTALL_DIR/properties` directory.
2. Adjust any of the following properties:

Option	Description
<code>purge_corePoolSize</code>	Minimum number of threads. The default is 1.
<code>purge_maxPoolSize</code>	Maximum number of threads. To scale External Purge for a higher load, increase the values for the <code>purge_maxPoolSize</code> and <code>purge_batchSize</code> properties. The default is 30.
<code>purge_keepAliveTime</code>	The External Purge JVM uses a pool of threads to distribute work. This value is the time in milliseconds that the threads will be kept alive if they become idle. If the workload is sporadic, you might need to use a larger value to keep the threads alive longer. The default is 10000.
<code>purge_maxQueueDepth</code>	The maximum depth of the in-memory queue. The maximum number of threads consumes a full batch each from this queue. The value must be kept sufficiently high so that enough work exists there for all the threads. A deep queue means fewer round-trips to the database for finding purge candidates. The default is 200000.
<code>purge_debug</code>	Indicates if the debug mode is active. The default is false.
<code>purge_scavengerSleepTime</code>	Amount of time the purge candidate scavenger sleeps between polls, in milliseconds. The default is 10000.
<code>purge_batchMode</code>	Indicates if you must wait to commit before reaching the batch count defined by the <code>purge_batchsize</code> parameter. Valid values: <ul style="list-style-type: none"> • <code>true</code> - (default) Wait to commit before reaching the batch count. You can commit before reaching the batch count if the depth of the queue is 0. • <code>false</code> - Do not wait to commit before reaching the batch count.
<code>purge_batchSize</code>	Batch size. The number of purge candidates per thread before a commit happens. The default is 100.

Option	Description
purge_dbPool	A new database pool that is used by just External Purge. When you create this pool, it references the <code>purge_maxPoolSize</code> property to create the number of threads needed for External Purge. When this pool is used, the database pool specified in Sterling B2B Integrator for all operations is not used in External Purge.
Note: The following parameters are available in V5.2.5_5 and higher	
purge_taskListCapacity	Can be used to set the capacity of the deferred delete file requests.
purge_numberDeleteThreads	Can be used to set the number of threads allocated for processing deferred delete file requests.

bin Directory Files Used by External Purge

You use a series of bin Directory files in the external purge.

File	Description
<code>control_extpurge.sh</code>	Starts and stops External Purge in UNIX and Linux. Valid arguments: <ul style="list-style-type: none"> • <code>start</code> - Starts External Purge. • <code>stop</code> - Stops External Purge. Runs <code>StopExternalPurge.sh</code>.
<code>ExternalPurgeConsole.sh/cmd</code>	For debugging. Starts External Purge in UNIX, Linux, or Windows and displays its progress in a console window. To stop the process, use Ctrl + c .
<code>InstallExternalPurgeWindowsService.cmd</code>	Installs External Purge in Windows. Run this command before running <code>startExternalPurgeWindowsService.cmd</code> .
<code>startExternalPurgeWindowsService.cmd</code>	Starts External Purge in Windows.
<code>StopExternalPurge.sh</code>	See description for <code>control_extpurge.sh</code> .
<code>stopExternalPurgeWindowsService.cmd</code>	Stops External Purge in Windows.
<code>uninstallExternalPurgeWindowsService.cmd</code>	Un-installs External Purge in Windows. Run this command after running <code>stopExternalPurgeWindowsService.cmd</code> .

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those

websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to

IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as shown in the next column.

© 2015.

Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. 2015.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Linear Tape-Open, LTO, the LTO Logo, Ultrium and the Ultrium Logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Connect Control Center[®], Connect:Direct[®], Connect:Enterprise[®], Gentran[®], Gentran[®]:Basic[®], Gentran:Control[®], Gentran:Director[®], Gentran:Plus[®], Gentran:Realtime[®], Gentran:Server[®], Gentran:Viewpoint[®], Sterling Commerce[™], Sterling Information Broker[®], and Sterling Integrator[®] are trademarks or registered trademarks of Sterling Commerce[®], Inc., an IBM Company.

Other company, product, and service names may be trademarks or service marks of others.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.



Product Number:

Printed in USA