Introduction

This article describes how to compress the embedded Apache Derby database included with the Integration Server product. This is very useful as the Apache Derby database disk file allocation will continue to increase as more rows are added to tables. Once rows are deleted, Derby does not automatically reduce the allocated files regardless of how many rows are removed. The only way to reduce this disk allocation is to perform a compression using Derby tools.

Overview

The steps required to perform the task are described as follows:

1. Stop the Integration Server.
2. Run the compression script attached to this article. It performs the following tasks:
   - Display available disk space to backup the database.
   - Displays information and prompts the user to continue.
   - Makes a backup of the 'db' directory in case compression fails and the database is corrupted.
   - Connects to the database.
   - Executes a Derby stored procedure reponsible for the compression.
   - Disconnects from the database.
   - If the compression fails, it will restore the backed up 'db' directory.
3. Start the Integration Server, confirm the Integration Server startup is successful and there are no new database-related errors in the server.log.
4. If the compressed database is corrupted, please do the following:
   - The script will restore the backed up database automatically.
   - Confirm the Integration Server starts up.
   - Contact support. Please provide the following:
     - Integration Server's server.log
     - Console output from compression script
     - Compressed corrupted directory.

- derby.log from the current working directory used to launch the script.

5. Remove the tar file for the embedded database backup if desired.

Caveats

- The following versions are supported. The current Derby version for an Integration Server release can be verified by looking in the MANIFEST.MF file of `<installpath>/IntegrationServer/common/lib/ext/derby.jar`

Integration Server	Derby
1015oct2022	10.15.2
1011oct2021	10.14.2
107oct2020	10.14.2
105oct2020	10.14.2
103oct2020	10.14.2
101oct2020	10.14.0 or 10.14.2

- Only tables from the Integration Server's APP schema are eligible for compression.
- There is no Windows .bat file for this tool. It's unconfirmed, but Linux script should work using Windows Subsystem for Linux, Cygwin etc.

Prerequsites

The Apache Derby tools jar is required to run this script. The jar can be downloaded from Maven Central. Once downloaded, please copy it to the same directory as the compression script.

Derby Version	Download URLs
10.15.2	derbytools-10.15.2.jar derby-10.15.2
10.14.2	derbytools-10.14.2.jar derby-10.14.2

Detailed Instructions

Unattended Execution

It's possible to run this script as an scheduled automated task via an environment variable.

For example…

```
export DERBY_UNATTENDED_COMPRESS="Enabled by John Doe"
```

The environment variable may be set to any value. As long as it is not null, the script does not prompt to continue; otherwise, the behavior is identical.

Stop the Integration Server

Only one process may have an embedded database open. The compression script creates a backup copy of the database prior to opening it.

Run the compression script

The script has the following command line parameters.

parameter	default	required?	description
-d	None	Yes	Absolute directory path to the Integration Server installation. Symbolic links are not supported.
-i	default	No	The instance name from *{installPath}/IntegrationServer/instances/{instance name}*
-t	IS_MONITOR, IS_SERVICE_STATS, IS_SERVER_STATS, IS_SERVER_DETAILS	No	A comma-delimited list of table names to compress
-ij	None	No	Launch the ij interactive SQL client. (The compression task does not execute.)
-h	None	No	Help message

The script creates a backup of the existing embedded db directory in the same directory. If the compression is successful, the backup is written to a compressed tar file and the directory is deleted. It's name pattern is *db_$(date '+%Y%m%d_%H%M%S').tar*.

The db directory backup tar files are intentionally not deleted by this script. Until the Integration Server is restarted and verified, backups are retained. Once the server is successfully started, the backup tar file may be deleted if desired.

If the Derby compression stored procedure returns a non-zero exit code, it is assumed the compression could have failed. In this case, the backup directory is restored automatically. The corrupted db directory is written to a tar file and the directory is removed. It's name pattern is *db-failed-compression_$(date '+%Y%m%d_%H%M%S').tar*.

The tar files created from a failed compression attempt are not deleted by the script. These files may be deleted whenver desired; however, the tar file could be used for diagnosing the cause of the failure.

Before and after the compression executes, a Derby diagnostic command executes. This command lists all the table names in the Integration Server database with disk space details including the estimated saving for compressing each table. Once this output is created, the script exits.

Script Validations

The script performs the following validation tasks:

- Confirms an installation directory is provided and it exists in the file system.
- Confirms the provided instance name contains a 'db' directory.
- Exits if the Integration Server is running as a service (If *.pid file exists)
- Confirms the Integration Server's Java installation exists. This Java installation is used to execute the Derby tools.
- Disables OS traps to prevent a script abort via Ctrl-C while the compression executes.
- An error message is produced if any of the provided table names do not exist in the schema. This is not a fatal error, so all of the listed tables will be processed.

For example: > ij> ERROR 38000: The exception 'java.sql.SQLSyntaxErrorException: 'ALTER TABLE' cannot be performed on 'APP.IS_SERVER_DETAILS' because it does not exist.' was thrown while evaluating an expression. ERROR 42Y55: 'ALTER TABLE' cannot be performed on 'APP.IS_SERVER_DETAILS' because it does not exist.

derby.log

The derby.log contains debugging information from the last script run. It is written to the current working directory used to launch the script.

Example Output

All script output is written to stdout (nothing to stderr). In this example, all of the rows from IS_MONITOR table were deleted to show a dramatic compression example. After compression, Derby expects an estimated savings of ESTIMSPACESAVING = 207880192 (pages). In this example, the db directory size was reduced from 865 megabytes to 16!

```
#### The embedded database is very large ####

$ du -chs db
```

```
1.9G    db
1.9G    total
$

#### After deleting 2.6 million rows, the directory is still very large ####

$ du -chs db
864M    db
864M    total
$

#### Launching the script ####

$ $rel/../110oct2022/tech-article-compress-derby/derby-compression.sh -d
/opt/data/installs/107oct2020


Welcome to the Apache Derby embedded database compression script.

This utility must be run while the Integration Server is running.
The required Apache derby jars must be present in the same directory as this
script.
The necessary instructions are included in the Tech Community article with
this script.

This script creates a backup of the existing embedded database prior to
compression.
It does not remove the backup when compression completes.
Verify the Integration Server start up is successful and there are no
unexpected database
    errors in the server.log before removing the backup copy.

The Derby diagnostic command shows disk space allocation for all tables in
the schema.
These details are provided before and after the compression to determine
which tables
    produced the most benefit.



    IS installation directory: /opt/data/installs/107oct2020
                IS instance: default
             Tables to compress:
IS_MONITOR,IS_SERVICE_STATS,IS_SERVER_STATS,IS_SERVER_DETAILS
         db directory and size: 849M
/opt/data/installs/107oct2020/IntegrationServer/instances/default/db

    Available space for db backup.
Filesystem      Size  Used Avail Use% Mounted on
```

```
/dev/sdb1       196G  148G   39G  80% /opt/data

Continue? (y/n) <n> y


Creating db backup:
/opt/data/installs/107oct2020/IntegrationServer/instances/default/db_20220910
_131023

Backup successful.
If compression fails, restore the db directory using this backup.

Creating /media/sf_rel/1015oct2022/../110oct2022/tech-article-compress-
derby/derby-run.sql ...

--------------------------

select count(*) from IS_MONITOR;
IS_SERVICE_STATS
select count(*) from IS_SERVICE_STATS;
IS_SERVER_STATS
select count(*) from IS_SERVER_STATS;
IS_SERVER_DETAILS
select count(*) from IS_SERVER_DETAILS;

SHOW SCHEMAS;
SET SCHEMA APP;
elapsedtime on;
SELECT CONGLOMERATENAME, ESTIMSPACESAVING, NUMALLOCATEDPAGES, NUMFREEPAGES
FROM SYS.SYSTABLES systabs,         TABLE (SYSCS_DIAG.SPACE_TABLE(
systabs.tablename )) AS T2         WHERE systabs.tabletype = 'T' and
ISINDEX=0 ORDER BY CONGLOMERATENAME;
EXIT;
--------------------------

/opt/data/installs/103oct2018/jvm/jvm/bin/java -cp
/media/sf_rel/1015oct2022/../110oct2022/tech-article-compress-derby/derby-
10.14.2.0.jar:/media/sf_rel/1015oct2022/../110oct2022/tech-article-compress-
derby/derbytools-10.14.2.0.jar -Dij.protocol="jdbc:derby:"    -
Dij.connection.IS_default="/opt/data/installs/103oct2018/IntegrationServer/in
stances/default/db/embedded;create=false;databaseName=APP"    -
Dderby.storage.indexStats.log=true\;    -Dderby.storage.indexStats.trace=log
-Dderby.storage.indexStats.auto=true    -Dderby.language.logQueryPlan=true
-Dderby.stream.error.logSeverityLevel=0    -Dderby.locks.deadlockTrace=true
-Dderby.locks.waitTimeout=300    -Dderby.debug.true=LogTrace
org.apache.derby.tools.ij <  /media/sf_rel/1015oct2022/../110oct2022/tech-
article-compress-derby/derby-run.sql
Executing SQL script.
```

```
ij version 10.14
IS_DEFAULT* -
jdbc:derby:/opt/data/installs/103oct2018/IntegrationServer/instances/default/
db/embedded
* = current connection
ij> TABLE_SCHEM
------------------------------
APP
NULLID
SQLJ
SYS
SYSCAT
SYSCS_DIAG
SYSCS_UTIL
SYSFUN
SYSIBM
SYSPROC
SYSSTAT

11 rows selected
ij> 0 rows inserted/updated/deleted
ij> ij> CONGLOMERATENAME
|ESTIMSPACESAVING    |NUMALLOCATEDPAGES   |NUMFREEPAGES
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
-----------------------------------
ATC_XREF
|0                   |1                   |0
COMPONENT_EVENT
|0                   |2                   |0
IS_ACCOUNT_LOCKING_DETAILS
|0                   |1                   |0
IS_ALERT
|0                   |1                   |0
IS_CERTIFICATE_MAP
|0                   |1                   |0
IS_CONSUMED_ALERTS
|0                   |1                   |0
IS_DATASTORE
|0                   |1                   |0
IS_GD_LOCKS
|0                   |1                   |0
IS_GD_TRANSACTIONS_IN
|0                   |1                   |0
IS_GD_TRANSACTIONS_OUT
|0                   |1                   |0
IS_KV_STORE
|0                   |1                   |0
IS_MONITOR
|207880192           |1                   |6344
```

IS_OAUTH_ACCESSTOKEN
|0 |1 |0
IS_OAUTH_CLIENTS
|0 |1 |0
IS_OAUTH_CLIENT_REDIR
|0 |1 |0
IS_OAUTH_CLIENT_SCOPE
|0 |1 |0
IS_OAUTH_REFRESHTOKEN
|0 |1 |0
IS_OAUTH_SCOPE
|0 |1 |0
IS_OAUTH_TOKEN
|0 |1 |0
IS_OAUTH_TOKEN_SCOPE
|0 |1 |0
IS_SERVER_DETAILS
|4096 |1 |1
IS_SERVER_STATS
|32768 |1 |8
IS_SERVICE_STATS
|325189632 |1 |79392
IS_TRIGGER_DOC
|0 |1 |0
IS_TRIGGER_STATE
|0 |1 |0
IS_USER_TASKS
|0 |1 |0
IS_WSRM_INVOKER_CLIENT
|0 |1 |0
IS_WSRM_INVOKER_SERVER
|0 |1 |0
IS_WSRM_MSGCTX_CLIENT
|0 |1 |0
IS_WSRM_MSGCTX_SERVER
|0 |1 |0
IS_WSRM_RMD_CLIENT
|0 |1 |0
IS_WSRM_RMD_SERVER
|0 |1 |0
IS_WSRM_RMS_CLIENT
|0 |1 |0
IS_WSRM_RMS_SERVER
|0 |1 |0
IS_WSRM_SENDER_CLIENT
|0 |1 |0
IS_WSRM_SENDER_SERVER
|0 |1 |0
WM_EXCL_DIST_LOCK
|0 |1 |0

```
WM_SHAR_DIST_LOCK
|0                  |1                  |0

38 rows selected
ELAPSED TIME = 131 milliseconds
ij> Elapsed time = 0 milliseconds


Compression completed.

Creating
/opt/data/installs/103oct2018/IntegrationServer/instances/default/db_20220910
_131023.tar. This may take a while.

tar: Removing leading `/' from member names


exit status = 0
Execution completed.
$


#### Checking disk space again ####

$ du -chs db
14M db
14M total
$
```

Additional Resources

Resource	URL	Description
ij documentation	User Guide	Documentation for the Derby ij command line SQL client
Admin Tools	Reference Guide	Detailed information for Derby's system admin tools. (Click the link for the required version and proceed to the 'refs' link.)