

# Backing up a derby database

This is a fragment page about how to backup a Derby database. Note that the database must not already be booted by another JVM, so this is NOT suitable for doing a live backup when Derby is running in Tomcat.

## Online backup

```
java org.apache.derby.tools.ij
ij>connect 'jdbc:derby:/path/to/your/db/sakaiadminx';
ij>CALL SYSCS_UTIL.SYSCS_BACKUP_DATABASE('/path/to/your/backup/');
```

While the backup is in progress, the execution of statements that require writes to the disk like insert, delete, and updates are blocked. Statements that require reads from the disk only, like select, proceed unblocked. Once the backup is complete, the system will let all blocked statements complete their execution. Online backup may decrease system throughput and user response time, so when possible schedule backups when update activity is low.

The backup process creates a directory with the same name as the database at the specified backup location. This directory contains all the files/directories that are in the database directory. The database directory typically has a data directory (seg0) to store all the data files for the tables and indexes, a transaction log directory (log) to store all the transaction log files, and a service.properties file that contains information to boot the database.

## Offline backup

The entire database directory can be copied when the system is offline using operating system commands. If the transaction log directory is different from the database directory, make sure that the transaction log directory is also copied. During the restore operation, transaction logs are necessary to bring the database to a consistent state.

Offline backup can also be performed without bringing the database system offline by freezing the system (blocking all write operations) while copying the database directory. The system procedure to do that is SYSCS\_UTIL.SYSCS\_FREEZE\_DATABASE(). Once the copy is complete, unfreeze (allow write operations) the system using the SYSCS\_UTIL.SYSCS\_UNFREEZE\_DATABASE() procedure; this procedure call will set the system back to normal state.

## Restore from a backup

The online/offline backup image of a database is used to restore the database to the state it was in, when the backup image was taken. The database is restored from the backup by specifying the connection URL attribute restoreFrom=<backup path> on the first connection to the database. Note that the backup path must include the database name in the backup, not just the backup location. The database will be restored to the location specified on the connection URL or to the current derby.system.home location if no path is specified on the connection URL. If a database with the same name exists at the current database location it will be removed first, and then restored from the backup.

```
ij>connect 'jdbc:derby:/path/to/your/db/sakaiadminx;restoreFrom=/path/to/your/backup/sakaiadminx';
```

Reference: <http://www.ibm.com/developerworks/db2/library/techarticle/dm-0502thalamati/>